

Optimal Inference After Model Selection

William Fithian* Dennis Sun Jonathan Taylor

Stanford University, Department of Statistics

May 1, 2015

Abstract

To perform inference after model selection, we propose controlling the *selective type I error*; i.e., the error rate of a test given that it was performed. By doing so, we recover long-run frequency properties among selected hypotheses analogous to those that apply in the classical (non-adaptive) context. Our proposal is closely related to data splitting and has a similar intuitive justification, but is more powerful. Exploiting the classical theory of Lehmann and Scheffé (1955), we derive most powerful unbiased selective tests and confidence intervals for inference in exponential family models after arbitrary selection procedures. For linear regression, we derive new selective z -tests that generalize recent proposals for inference after model selection and improve on their power, and new selective t -tests that do not require knowledge of the error variance σ^2 .

1 Introduction

A typical statistical investigation can be thought of as consisting of two stages:

- 1. Selection:** The analyst chooses a probabilistic model for the data at hand, and formulates testing, estimation, or other problems in terms of unknown aspects of that model.
- 2. Inference:** The analyst attempts the chosen problems using the data and the selected model.

Informally, the selection stage determines what questions to ask, and the inference stage answers those questions. Most statistical methods carry an implicit assumption that selection is *non-adaptive* — that is, choices about which model to use, hypothesis to test, or parameter to estimate, are made before seeing the data. *Adaptive selection* (also known colloquially as “data snooping”) violates this assumption, formally invalidating any subsequent inference.

In some cases, it is possible to specify the question prior to collecting the data—for instance, if the data are governed by some known physical law. However, in most applications, the choice of question is at least partially guided by the data. For example, we often perform exploratory analyses to decide which predictors or interactions to include in a regression model or to check whether the assumptions of a test are satisfied. In this work, we define a framework for valid inference after adaptive selection.

If we do not account properly for adaptive model selection, the resulting inferences can have troubling frequency properties, as we now illustrate with an example.

*To whom correspondence should be addressed

Example 1 (File Drawer Effect). Suppose a scientist observes n independent observations $Y_i \sim N(\mu_i, 1)$. He focuses only on the apparently large effects, selecting only the indices i for which $|Y_i| > 1$, i.e.

$$\hat{I} = \{i : |Y_i| > 1\}.$$

He wishes to test $H_{0,i} : \mu_i = 0$ for each $i \in \hat{I}$ at the $\alpha = 0.05$ significance level. Most scientists intuitively recognize that the nominal test that rejects $H_{0,i}$ when $|Y_i| > 1.96$ is invalidated by the selection.

What exactly is “invalid” about this test? After all, the probability of falsely rejecting a given $H_{0,i}$ is still α , since most of the time, $H_{0,i}$ is simply not tested at all. Rather, the troubling feature is that the error rate among the hypotheses *selected* for testing is possibly much higher than α . To be precise, let n_0 be the number of true null effects and suppose $n_0 \rightarrow \infty$ as $n \rightarrow \infty$. Then, in the long run, the fraction of errors among the true nulls we test is

$$\begin{aligned} \frac{\# \text{ false rejections}}{\# \text{ true nulls selected}} &= \frac{\frac{1}{n_0} \sum_{i: H_{0,i} \text{ true}} 1\{i \in \hat{I}, \text{ reject } H_{0,i}\}}{\frac{1}{n_0} \sum_{i: H_{0,i} \text{ true}} 1\{i \in \hat{I}\}} \\ &\rightarrow \frac{\mathbb{P}_{H_{0,i}}(i \in \hat{I}, \text{ reject } H_{0,i})}{\mathbb{P}(i \in \hat{I})} \\ &= \mathbb{P}_{H_{0,i}}(\text{reject } H_{0,i} \mid i \in \hat{I}), \end{aligned} \tag{1}$$

which for the nominal test is $\Phi(-1.96)/\Phi(-1) \approx .16$.

Thus, we see that (1), the probability of a false rejection conditional on selection, is a natural error criterion to control in the presence of selection. In this example, we can directly control (1) at level $\alpha = 0.05$ simply by finding the critical value c solving

$$\mathbb{P}_{H_{0,i}}(|Y_i| > c \mid |Y_i| > 1) = 0.05.$$

In this case $c = 2.41$, which is more stringent than the nominal 1.96 cutoff.

This paper will develop a theory for inference after selection based on controlling the *selective type I error rate* (1). Our guiding principle is:

The answer must be valid, given that the question was asked.

For all its disarming simplicity, Example 1 can be regarded as a stylized model of science. Imagine that each Y_i represents an estimated effect size from a scientific study. However, only the large estimates are ever published—a caricature which may not be too far from the truth, as recently demonstrated by Franco et al. (2014). To compound the problem, there may be many reasonable methodologies to choose from, even once the analyst has decided roughly what scientific question to address (Gelman and Loken, 2013). Because of the resulting selection bias, the error rate among published claims may be very high, leading even to speculation that “most published research findings are false” (Ioannidis, 2005). Thus, selection effects may be a partial explanation for the replicability crisis and decline effect reported in the scientific community (Yong, 2012) and the popular media (Johnson, 2014).

The setting of Example 1 has been studied extensively in the literature of simultaneous and selective inference, and several authors have proposed adjusting for selection by means of conditional inference. Zöllner and Pritchard (2007) and Zhong and Prentice (2008) construct selection-adjusted estimators and intervals for genome-wide association studies for genes that pass a fixed initial significance threshold, based on a conditional Gaussian likelihood. Cohen and Sackrowitz

(1989) obtain unbiased estimates for the mean of the population whose sample mean is largest by conditioning on the ordering of the observed sample means, and Sampson and Sill (2005) and Sill and Sampson (2009) apply the same idea to obtain estimates for the best-performing drug in an adaptive clinical trial design. Hedges (1984) and Hedges (1992) propose methods to adjust for the file drawer effect in meta-analysis when scientists only publish significant results.

Another framework for selection adjustment is proposed by Benjamini and Yekutieli (2005), who consider the problem of constructing intervals for a number R of parameters selected after viewing the data. Letting V denote the number of non-covering intervals among those constructed, they define the *false coverage-statement rate* (FCR) as the expected fraction $V/\max(R, 1)$ of non-covering intervals. Controlling the FCR at level α thus amounts to “coverage on the average, among selected intervals.” As we will see further in Section 8, FCR control is closely related to the selective error control criterion we propose. In fact, Weinstein et al. (2013) employ conditional inference to construct FCR-controlling intervals in the context of Example 1. Rosenblatt and Benjamini (2014) propose a similar method for finding correlated regions of the brain, also with a view toward FCR control.

1.1 Conditioning on Selection

In classical statistical inference, the notion of “inference after selection” does not exist. The analyst must specify the model, as well as the hypothesis to be tested, in advance of looking at the data. A classical level- α test for a hypothesis H_0 under model M must control the usual or *nominal type I error rate*:

$$\mathbb{P}_{M,H_0}(\text{reject } H_0) \leq \alpha. \quad (2)$$

The subscript in (2) reminds us that the probability is computed under the assumption that the data Y are generated from model M , and H_0 is true; if M is misspecified, there are no guarantees on the rejection probability.

While ruling out model selection avoids the selection problem altogether, it does not realistically describe most statistical practice: statisticians are trained to check their models and to tweak them if they diagnose a problem. Under the purist view, model checking is technically forbidden, since it leaves open the possibility that the model will change after we see the data. We will argue that if the model and hypothesis are selected randomly, we should instead control the selective type I error rate

$$\mathbb{P}_{M,H_0}(\text{reject } H_0 \mid (M, H_0) \text{ selected}) \leq \alpha. \quad (3)$$

One can argue that models and hypotheses are practically never truly fixed but are chosen randomly, since they are based on the outcomes of previous experiments in the (random) scientific process. Typically, we ignore the random selection and use classical tests that control (2), implicitly assuming that the randomness in selecting M and H_0 is independent of the data used for inference. In that case,

$$\mathbb{P}_{M,H_0}(\text{reject } H_0 \mid (M, H_0) \text{ selected}) = \mathbb{P}_{M,H_0}(\text{reject } H_0). \quad (4)$$

It may seem pedantic to point out that model selection is random if based on previous experiments. Even so, this viewpoint gives us a prescription for what to do when science does not dictate a model. If it is possible to split the data $Y = (Y_1, Y_2)$ with Y_1 independent of Y_2 , then we can imitate the scientific process by setting aside Y_1 for selection and Y_2 for inference. If selection depends on Y_1 only, then any nominal level- α test based on the value of Y_2 will satisfy (4), so the nominal test based on Y_2 also controls the selective error (3).

This meta-algorithm for generating selective procedures from nominal ones is called *data splitting* or *sample splitting*. The idea dates back at least as far as Cox (1975), and, despite the paucity of literature on the topic, is common wisdom among practitioners. For example, it is customary

in genetics to use one cohort to identify loci of interest and a separate cohort to confirm them (Sladek et al., 2007).

The popularity of data splitting owes in no small part to its transparent justification, which non-experts can easily appreciate: if we imagine that Y_1 is observed “first,” then we can proceed to analyze Y_2 as though model selection took place “ahead of time.” Equation (4) guarantees that this temporal metaphor will not lead us astray even if it does not describe how Y_1 and Y_2 were actually collected.

Data splitting elegantly solves the problem of controlling selective error, but at a cost. It not only reduces the amount of data available for inference, but also reduces the amount of data available for selection. Furthermore, it is not always possible; for example, spatial and time series data often exhibit autocorrelation that rules out splitting the data into independent parts.

In this article, we propose directly controlling the selective error rate (3) by conditioning on the event that (M, H_0) is selected. As with data splitting, we treat the data as though it were revealed in stages: in the first stage, we “observe” just enough data to resolve the decision of whether to test (M, H_0) , after which we can treat the data $(Y \mid (M, H_0) \text{ selected})$ as “not yet observed” when stage two commences.

The intuition of the above paragraph can be expressed formally in terms of the filtration

$$\mathcal{F}_0 \quad \underbrace{\subseteq}_{\text{used for selection}} \quad \mathcal{F}(\mathbf{1}_A(Y)) \quad \underbrace{\subseteq}_{\text{used for inference}} \quad \mathcal{F}(Y), \quad (5)$$

where $\mathcal{F}(Z)$ denotes the σ -algebra generated by random variable Z (informally, everything we know about the data after observing Z), \mathcal{F}_0 is the trivial σ -algebra (representing complete ignorance), and A is the *selection event* $\{(M, H_0) \text{ selected}\}$. We can think of “time” as progressing from left to right in (5). In stage one, we learn just enough to decide whether to test (M, H_0) , and no more, advancing our state of knowledge from \mathcal{F}_0 to $\mathcal{F}(\mathbf{1}_A(Y))$. We then begin stage two, in which we discover the actual value of Y , advancing our knowledge to $\mathcal{F}(Y)$. Because our selection decision is made at the end of stage one, everything revealed during stage two is fair game for inference.

In effect, controlling the type I error conditional on A prevents us from appealing to the fact that $Y \in A$ as evidence against H_0 . Even if $Y \in A$ is extremely surprising under H_0 , we still will not reject unless we are surprised anew in the second stage. In this sense, conditioning on a random variable discards the information it carries about any parameter or hypothesis of interest. By contrast with data splitting, which can be viewed as conditioning on Y_1 instead of $\mathbf{1}_A(Y_1)$, we advocate discarding as little data as possible and reserving the rest for stage two. This frugality results in a more efficient division of the information carried by Y — *data carving*, to introduce an evocative metaphor.

1.2 Outline

In Section 2 we formalize the problem of selective inference, discuss general properties of selective error control, and address key conceptual questions. Conditioning on the selection event effectively discards the information used for selection, but some information is left over for second-stage inference. We will also see that a major advantage of selective error control is that it allows us to consider only one model at a time when designing tests and intervals, even if *a priori* there are many models under consideration.

If $\mathcal{L}(Y)$, the law of random variable Y , follows an exponential family model, then for any event A , $\mathcal{L}(Y \mid A)$ follows a closely related exponential family model. As a result, selective inference dovetails naturally with the classical optimality theory of Lehmann and Scheffé (1955); Section 3 briefly reviews this theory and derives most powerful unbiased selective tests in arbitrary exponential family models after arbitrary model selection procedures. Because conditioning on more data

than is necessary saps the power of second-stage tests, data splitting yields inadmissible selective tests under general conditions.

Section 5 gives some general strategies for computing rejection cutoffs for the tests prescribed in Section 3, while Sections 4–6 derive selective tests in specific examples. Section 4 focuses on the case of linear regression, generalizing the recent proposals of Taylor et al. (2014), Lee et al. (2013), and others. We derive new, more powerful selective z -tests, as well as t -tests that do not require knowledge of the error variance σ^2 .

Several simulations in Section 7 compare the post-lasso selective z -test with data splitting, and illustrate a *selection–inference tradeoff*, between using more data in the initial stage and reserving more information for the second stage. Section 8 compares and contrasts selective inference with multiple inference, and Section 9 concludes.

2 The Problem of Selective Inference

2.1 Example: Regression and the Lasso

In the previous section, we motivated the idea of conditioning on selection. Arguably, the most familiar example of this “selection” is variable selection in linear regression. In regression, the observed data $Y \in \mathbb{R}^n$ is assumed to be generated from a multivariate normal distribution

$$Y \sim N_n(\mu, \sigma^2 I_n). \quad (6)$$

The goal is to model the mean μ as a linear function of predictors X_j , $j = 1, \dots, p$. To obtain a more parsimonious model (or simply an identifiable model when $p > n$), researchers will often use only a subset $M \subseteq \{1, \dots, p\}$ of the predictors. Each subset M leads to a different probabilistic model corresponding to the assumption $\mu = X_M \beta^M$, where X_M denotes the matrix consisting of columns X_j for $j \in M$. Then, it is customary to report tests of $H_{0,j}^M : \beta_j^M = 0$ for each coefficient in the model. If M was chosen in a data-dependent way, then to control selective error we must condition on having selected $(M, H_{0,j}^M)$, which in this case is the same as conditioning on having selected model M .

There are many data-driven methods for variable selection in linear regression, ranging from AIC minimization to forward stepwise selection, cf. Hastie et al. (2009). We will consider one procedure in particular, based on the lasso, mostly because selective inference in the context of the lasso (Lee et al., 2013) was a main motivation for the present work. The lasso (Tibshirani, 1996) provides an estimate of $\beta \in \mathbb{R}^p$ that solves

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \|Y - X\beta\|_2^2 + \lambda \|\beta\|_1, \quad (7)$$

where X is the “full” matrix consisting of all p predictors. The first term is the usual least-squares objective, while the second term encourages many of the coefficients to be exactly zero. Because of this property, it makes sense to define the model “selected” by the lasso to be the set of variables with non-zero coefficients, i.e.,

$$\widehat{M}(Y) = \{j : \hat{\beta}_j \neq 0\}.$$

Notice that $\widehat{M}(Y)$ can take on up to 2^p possible values, one for each subset of $\{1, \dots, p\}$. The regions $A_M = \{y : \widehat{M}(y) = M\}$ form a partition of \mathbb{R}^n into regions that correspond to each model. To control the selective error after selecting a particular M , we must condition on the event that Y landed in A_M . The partition for a lasso problem with $p = 3$ variables in $n = 2$ dimensions is shown in Figure 1. An explicit characterization of the lasso partition can be found in Lee et al. (2013); see also Harris (2014) for more plots visualizing the way the lasso partitions the sample space. If we used a different selection procedure, we would obtain a different partition. Characterizations




Figure 1: An example of the lasso with $n = 2$ observations and $p = 3$ variables. We base tests on the distribution of Y , conditional on its landing in the highlighted region.

of the partitions in forward stepwise selection and marginal screening can be found in Loftus and Taylor (2014) and Lee and Taylor (2014), respectively.

Let us imagine that in stage one, we loaded the data into a software package and computed $\widehat{M}(Y)$, but we remain otherwise ignorant of the value Y — that is, we have observed *which* of the regions Y falls into but not *where* Y is in that region. Now that we have chosen the model, we will construct tests of $H_{0,j}^M : \beta_j^M = 0$ for each of the selected variables. In the example shown in Figure 1, we selected variables 1 and 3 and thus test the two hypotheses

$$\begin{aligned} H_{0,1}^{\{1,3\}} : \beta_1^{\{1,3\}} &= 0 \\ H_{0,3}^{\{1,3\}} : \beta_3^{\{1,3\}} &= 0. \end{aligned}$$

Notice that we have to be careful to always specify the model along with the coefficient, since the coefficient for variable j does not necessarily have a consistent interpretation across different models. Each regression coefficient summarizes the effect of that variable, adjusting for the other variables in the model. For example, “What is the effect of IQ on salary?” is a genuinely different question from “What is the effect of IQ on salary, after adjusting for years of education?” Both questions are meaningful, but they are fundamentally different.¹

Having chosen the model M and conditioned on the selection, we will base our tests on the precise location of Y , which we do not know yet. Conditionally, Y is not Gaussian, but it does follow an exponential family. As a result, we can appeal to the classical theory of Lehmann and Scheffé (1955) to construct tests or confidence intervals for its natural parameters, which are β^M if σ^2 is known, and otherwise are $(\beta^M/\sigma^2, 1/\sigma^2)$.

With this concrete example in mind, we will now develop a general framework of selective inference that is much more broadly applicable. Because we are explicitly allowing models and hypotheses to be random, it is necessary to carefully define our inferential goals. We first discuss selective inference in the context of hypothesis testing. The closely related developments for confidence intervals will follow in Section 2.3.

¹We use the word “effect” here informally to refer to a regression coefficient, recognizing that regression cannot establish causal claims on its own.

2.2 Selective Hypothesis Tests

We now introduce notation that we will use for the remainder of the article. Assume that our data Y lies in some measurable space $(\mathcal{Y}, \mathcal{F})$, with unknown sampling distribution $Y \sim F$. The analyst's task is to pose a reasonable probability model M — i.e., a family of distributions which she believes contains F — and then carry out inference based on the observation Y .

Let \mathcal{Q} denote the *question space* of inference problems q we might tackle. A hypothesis testing problem is a pair $q = (M, H_0)$ of a model M and null hypothesis H_0 , by which we mean a submodel $H_0 \subseteq M$.² We write $M(q)$ and $H_0(q)$ for the model and hypothesis corresponding to q . Without loss of generality, we assume $H_0(q)$ is tested against the alternative hypothesis $H_1(q) = M(q) \setminus H_0(q)$. To avoid measurability issues, we will assume throughout that \mathcal{Q} is countable, although our framework can be extended to other question spaces with additional care.

In Section 2.1 where we test each variable in a selected regression model, the question space is

$$\mathcal{Q} = \{(M, H_{0,j}^M) : M \subseteq \{1, \dots, p\}, j \in M\}.$$

Note our slight abuse of notation in using M interchangeably to refer both to a subset of variable indices and to the corresponding probability model $\{N_n(X_M \beta_M, \sigma^2 I_n) : \beta \in \mathbb{R}^{|M|}\}$.

We model selective inference as a process with two distinct stages:

1. **Selection:** From the collection \mathcal{Q} of possible questions, the analyst selects a subset $\widehat{\mathcal{Q}}(Y) \subseteq \mathcal{Q}$ to test, based on the data.
2. **Inference:** The analyst performs a hypothesis test of $H_0(q)$ against $M(q) \setminus H_0(q)$ for each $q \in \widehat{\mathcal{Q}}$.

In the case of the simple regression example shown in Figure 1, where we selected variables 1 and 3, $\widehat{\mathcal{Q}}$ would consist of the hypotheses for each of the two variables in the model. To be completely explicit,

$$\widehat{\mathcal{Q}}(Y) = \left\{ \left(\{1, 3\}, H_{0,1}^{\{1,3\}} \right), \left(\{1, 3\}, H_{0,3}^{\{1,3\}} \right) \right\}.$$

A correctly specified model M is one that contains the true sampling distribution F . Importantly, we expressly do not assume that all — or any — of the candidate models are correctly specified. Because the analyst must choose M without knowing F , she could choose poorly, in which case there may be no formal guarantees on the behavior of the test she performs in stage two. Some degree of misspecification is the rule rather than the exception in most real statistical applications, whether models are specified adaptively or non-adaptively. Our analyst would be in the same position if she were to select a (probably wrong) model using Y , then use that model to perform a test on new data Y^* collected in a confirmatory experiment. See Section 2.5.2 for further discussion of this issue.

For our purposes, a *hypothesis test* is a function $\phi(y)$ taking values in $[0, 1]$, representing the probability of rejecting H_0 if $Y = y$. In most cases, the value of the function will be either 0 or 1, but with discrete variables, randomization may be necessary to achieve exact level α .

To adjust for selection in testing q , we condition on the event that the question was asked, which we describe by the selection event

$$A_q = \{q \in \widehat{\mathcal{Q}}(Y)\}, \tag{8}$$

i.e., the event that q is among the questions asked. In general, the selection events for different questions are not disjoint. In the regression example, we only ever test $H_{0,j}^M$ when model M is selected, so conditioning on A_q is equivalent to simply conditioning on \widehat{M} .

²We identify a “null hypothesis” like $H_0 : \mu(F) = 0$ with the corresponding subfamily or “null model” $\{F \in M : \mu(F) = 0\}$. This should remind us that the error guarantees of a test do not necessarily extend beyond the model it was designed for.

In selective inference, we are mainly interested in the properties of a test ϕ_q for a question q , conditional on A_q . We say that ϕ_q controls *selective type I error* at level α if

$$\mathbb{E}_F[\phi_q(Y) | A_q] \leq \alpha, \quad \text{for all } F \in H_0. \quad (9)$$

and define its *selective power function* as

$$\text{Pow}_{\phi_q}(F | A_q) = \mathbb{E}_F[\phi_q(Y) | A_q]. \quad (10)$$

Because \mathcal{Q} is countable, the only relevant q are those for which $\mathbb{P}(A_q) > 0$.

Notice that only the model M and hypothesis H_0 are relevant for defining the selective level of a test. This means that in designing valid ϕ_q , we can concentrate on one q at a time, even if there are many mutually incompatible candidate models in \mathcal{Q} . As long as each ϕ_q controls the selective error at level α given its selection event A_q , then a global error is also controlled:

$$\frac{\mathbb{E}[\# \text{ false rejections}]}{\mathbb{E}[\# \text{ true nulls selected}]} \leq \alpha, \quad (11)$$

provided that the denominator is finite. (11) holds for countable \mathcal{Q} regardless of the dependence structure across different q . The fact that we can design tests one q at a time makes it much easier to devise selective tests in concrete examples, which we take up in Sections 3–6.

Suppose that each scientist in a discipline controls the selective error rate for each of his or her own experiments. Then the discipline as a whole will achieve long-run control of the type I error rate among true *selected* null hypotheses, in the same sense as they would if there were no selection. No coordination is required between different research groups.

Proposition 1 (Discipline-Wide Error Control). *Suppose there are n independently operating research groups in a scientific discipline with a shared, countable question space \mathcal{Q} . Research group i collects data $Y_i \sim F_i$, applies selection rule $\widehat{\mathcal{Q}}_i(Y_i) \subseteq \mathcal{Q}$, and carries out selective level- α tests $(\phi_{q,i}(y_i), q \in \widehat{\mathcal{Q}}_i)$. Assume each research group has probability at least $\delta > 0$ of carrying out at least one test of a true null, and for some common $B < \infty$,*

$$\mathbb{E}_{F_i} \left[|\widehat{\mathcal{Q}}_i(Y_i)|^2 \right] \leq B, \quad \text{for all } i.$$

Then as n grows, the discipline as a whole achieves long-run control over the frequentist error rate

$$\limsup_{n \rightarrow \infty} \frac{\# \text{ false rejections}}{\# \text{ true nulls selected}} \stackrel{a.s.}{\leq} \alpha. \quad (12)$$

The proof is deferred to Appendix A. Though the assumption of independence across research groups can be weakened without necessarily affecting the conclusion, we do not pursue such generalizations. Note that there is no counterpart to Proposition 1 for multiple-inference error rates such as the false discovery rate (FDR) (Benjamini and Hochberg, 1995) or familywise error rate (FWER); even if every research group controls its own FWER or FDR at level α , there is no guarantee we will control FWER or FDR after aggregating across the different groups.

2.3 Selective Confidence Intervals

If the goal is instead to form confidence intervals for a parameter $\theta(F)$, it is more convenient to think of \mathcal{Q} as containing pairs $q = (M, \theta(\cdot))$ of a model and a parameter. By analogy to (9), we will call a set $C(Y)$ a $(1 - \alpha)$ *selective confidence set* if

$$\mathbb{P}_F(\theta(F) \in C(Y) | A_q) \geq 1 - \alpha, \quad \text{for all } F \in M. \quad (13)$$

The next result establishes that selective confidence sets can be obtained by inverting selective tests, as one would expect by analogy to the classical case.




Figure 2: Instead of conditioning on the selection event A_q that question q is asked, we can condition on a finer event, the value of the random variable S_q . We call S_q the *selection variable*.

Proposition 2 (Duality of Selective Tests and Confidence Sets). *Suppose we form a confidence interval for $\theta(F)$ on the event A_q . Suppose also that on this event, we form a test ϕ_t of $H_{0,t} = \{F : \theta(F) = t\}$ for all t . Let $C(Y)$ be the set of t for which ϕ_t does not (always) reject:*

$$C(Y) = \{t : \phi_t(Y) < 1\}. \quad (14)$$

If each ϕ_t is a selective level- α test, then $C(Y)$ is a selective $(1 - \alpha)$ confidence set.

Proof. The selective non-coverage probability is

$$\mathbb{P}_F(\theta(F) \notin C(Y) | A_q) = \mathbb{P}_F(\phi_{\theta(F)}(Y) = 1 | A_q) \leq \mathbb{E}_F[\phi_{\theta(F)}(Y) | A_q] \leq \alpha.$$

□

2.4 Conditioning Discards Information

Because performing inference conditional on a random variable effectively disqualifies that variable as evidence against a hypothesis, we will typically want to condition on as little data as possible in stage two. Even so, some selective inference procedures condition on more than A_q . For example, data splitting can be viewed as inference conditional on Y_1 , the part of the data used for selection. More generally, we say a *selection variable* is any variable $S_q(Y)$ whose level sets partition the sample space more finely than A_q ; i.e., $A_q \in \mathcal{F}(S_q)$. Informally, we can think of conditioning on a finer partition of A_q , as shown in Figure 2.

We say ϕ controls the *selective type I error with respect to S_q* at level α if the error rate is less than α given $S_q = s$ for $\{S_q = s\} \subseteq A_q$. More formally,

$$\mathbb{E}_F[\phi(Y)\mathbf{1}_{A_q}(Y) | S_q] \stackrel{\text{a.s.}}{\leq} \alpha, \quad \text{for all } F \in H_0 \quad (15)$$

Taking $S_q(y) = \mathbf{1}_{A_q}(y)$, the coarsest possible selection variable, recovers the baseline selective type I error in (9). The definition of a selective confidence set may be generalized in the same way.

Generalizing (5) to finer selection variables gives

$$\mathcal{F}_0 \quad \underbrace{\subseteq}_{\text{used for selection}} \quad \mathcal{F}(S(Y)) \quad \underbrace{\subseteq}_{\text{used for inference}} \quad \mathcal{F}(Y), \quad (16)$$

suggesting that the more we refine $S(Y)$, the less data we have left for second-stage inference. Indeed, the finer S is, the more stringent is the requirement (15):

Proposition 3 (Monotonicity of Selective Error). *Suppose $\mathcal{F}(S_1) \subseteq \mathcal{F}(S_2)$. If ϕ controls the type I error rate at level α for $q = (M, H_0)$ w.r.t. the finer selection variable S_2 , then it also controls the type I error rate at level α w.r.t. the coarser S_1 .*

Proof. If $F \in H_0$, then

$$\mathbb{E}_F[\phi(Y)\mathbf{1}_A(Y) | S_1] = \mathbb{E}_F[\mathbb{E}_F[\phi(Y)\mathbf{1}_A(Y) | S_2] | S_1] \stackrel{\text{a.s.}}{\leq} \alpha.$$

□

Because $S(y) = \mathbf{1}_A(y)$ is the coarsest possible choice, a test controlling the type I error w.r.t. any other selection variable also controls the selective error in (9). At the other extreme, if $S(y) = y$, then we cannot improve on the trivial “coin-flip” test $\phi(y) \equiv \alpha$. Proposition 3 suggests that we will typically sacrifice power as we move from coarser to finer selection variables. Even so, refining the selection variable can be useful for computational reasons. For example, in the case of the lasso, by conditioning additionally on the signs of the nonzero $\hat{\beta}_j$, the selection event becomes a convex region instead of up to $2^{|\bar{M}|}$ disjoint convex regions (Lee et al., 2013). Another valid reason to refine S_q beyond $\mathbf{1}_{A_q}$ is to strengthen our inferential guarantees in a meaningful way; for example, we can achieve false coverage-statement rate (FCR) control by choosing $S_q = (\mathbf{1}_{A_q}(Y), |\hat{Q}(Y)|)$ (see Section 8, Proposition 11).

Data splitting corresponds to setting every selection variable equal to $S = Y_1$. As a result, data splitting does not use all the information that remains after conditioning on A , as we see informally in the three-stage filtration

$$\mathcal{F}_0 \quad \underbrace{\subseteq}_{\text{used for selection}} \quad \mathcal{F}(\mathbf{1}_A(Y_1)) \quad \underbrace{\subseteq}_{\text{wasted}} \quad \mathcal{F}(Y_1) \quad \underbrace{\subseteq}_{\text{used for inference}} \quad \mathcal{F}(Y_1, Y_2). \quad (17)$$

As we will see in Section 3.2, this waste of information means that data splitting is inadmissible under fairly general conditions.

We can quantify the amount of leftover information in terms of the Fisher information that remains in the conditional law of Y given S . In a smooth parametric model, we can decompose the Hessian of the log-likelihood as

$$\nabla^2 \ell(\theta; Y) = \nabla^2 \ell(\theta; S) + \nabla^2 \ell(\theta; Y | S) \quad (18)$$

The conditional expectation

$$\mathcal{I}_{Y|S}(\theta; S) = -\mathbb{E}[\nabla^2 \ell(\theta; Y | S) | S] \quad (19)$$

is the *leftover Fisher information* after selection at $S(Y)$. Taking expectations in (18), we obtain

$$\mathbb{E}[\mathcal{I}_{Y|S}(\theta; S)] = \mathcal{I}_Y(\theta) - \mathcal{I}_S(\theta) \preceq \mathcal{I}_Y(\theta). \quad (20)$$

Thus, on average, the price of conditioning on S — the price of selection — is the information S carries about θ .³ In some cases this loss may be quite small, which a simple example elucidates.


Example 2. Consider selective inference under the univariate Gaussian model

$$Y \sim N(\mu, 1), \quad (21)$$

after conditioning on the selection event $A = \{Y > 3\}$.

Figure 3a plots the leftover information as a function of μ . If $\mu \ll 3$, there is very little information in the conditional distribution: whether $\mu = -10$ or $\mu = -11$, Y is conditionally highly concentrated on 3. By contrast, if $\mu \gg 3$, then $\mathbb{P}_\mu(A) \approx 1$, the conditional law is practically no different from the marginal law, and virtually no information is lost in the conditioning.

³Note that we do not necessarily have $\mathcal{I}_{Y|S}(\theta; S) \preceq \mathcal{I}_Y(\theta)$ for every S . In fact there are interesting counterexamples where $\mathcal{I}_{Y|A}(\theta) \gg \mathcal{I}_Y(\theta)$ for certain θ , but we will not take them up here.



(a) Leftover Fisher information as a function of μ . For $\mu \ll 3$, then there is very little information in the conditional distribution, since Y is conditionally highly concentrated on 3. For $\mu \gg 3$, then $\mathbb{P}_\mu(A) \approx 1$ and virtually no information is lost.

(b) Confidence intervals from inverting the UMPU tests of Section 3. For $Y \gg 3$, the interval essentially coincides with the nominal interval $Y \pm 1.96$. For Y close to 3, the wide interval reflects potentially severe selection bias.

Figure 3: Univariate Gaussian. $Y \sim N(\mu, 1)$ with selection event $A = \{Y > 3\}$.

Figure 3b shows the confidence intervals that result from inverting the tests described in Section 3. When $Y \gg 3$, the interval essentially coincides with the nominal interval $Y \pm 1.96$ because there is hardly any selection bias and no real adjustment is necessary. By contrast, when Y is close to 3 it is potentially subject to severe selection bias. This fact is reflected by the confidence interval, which is both longer than the nominal interval and centered at a value significantly less than Y .

2.5 Conceptual Questions

We now pause to address conceptual objections that may have occurred to a skeptical reader.

2.5.1 How Can the Model Be Random?

In our framework, inference is based on a probabilistic model that is allowed to be chosen randomly, based on the data Y ; the reader may wonder whether that randomness muddies the interpretation of whatever inference is carried out in the second stage.

First, note that in our framework the *true* sampling distribution F is not “selected” in any sense; it is entirely outside the analyst’s control. The only thing selected is the *working model*, a tool the analyst uses to carry out inference, which may or may not include the true F .

Thoughtful skeptics may find reasons for concern about any approach — data carving, data splitting, or selecting a model based on a previous experiment — in which a random model is selected, believing that statistical testing is only appropriate when a probabilistic model can be based purely on convincing theoretical considerations. We answer only that this point of view would rule out most scientific inquiries for which statistics is ever used. However, if one is comfortable with choosing a random model based on data splitting or a previous experiment, we see no special reason to be any more concerned about choosing a random model based on data carving.

In any case, it is by no means required that the model M be random. For example, in our clinical trial example of Section 6.1, the probabilistic model is always the same but we choose

which null hypotheses to test after inspecting the data. The same is true of the saturated-model selective z -test described in 4.2.

2.5.2 What if the Selected Model is Wrong?

If we were not writing about model selection, we might have begun by stating a formal mathematical assumption that the sampling distribution F belongs to a known model M , and then devised a test ϕ that behaves well when $F \in M$. The same ϕ might not work well at all for $F \notin M$: for example, if we choose to apply the one-sample t -test of $\mu = 0$ to a sample Y_1, \dots, Y_n whose observations are highly correlated, then the probability of rejection may be a great deal larger than the nominal α , even if $\mathbb{E}[Y_i] = 0$. This is not a mistake in the formal theory, nor does it make the t -test an inherently invalid test; rather, the validity or invalidity of a test is defined with respect to its behavior when $F \in H_0 \subseteq M$.

In any given application, the analyst must choose from among many statistical methods knowing that each one is designed to work under a particular set of assumptions about F — i.e., under a particular model M . Because our theory encompasses both the choice and the subsequent analysis, it would not be sensible to assume that the analyst is infallible and always selects a correct model. Typically some candidate models M are correctly specified (i.e., $F \in M$), others are not ($F \notin M$), and the analyst can never know for sure which are which.

Of course, the possibility of misspecification is not restricted to adaptive procedures like data carving and data splitting: selecting an inappropriate model *after* seeing the data leaves us no better or worse off than if we had chosen the same inappropriate model *before* seeing the data. Each ϕ_q or C_q is designed with respect to a particular model $M(q)$, and properties like selective type I error control or selective coverage only constrain its behavior for $F \in M(q)$.

There is a separate question of robustness: if $F \notin M$ but is “close” in some sense, we may still want our procedure to behave predictably. However, even if some model gives a reasonable approximation to $\mathcal{L}(Y)$, there is no guarantee that the induced model for $\mathcal{L}(Y | A)$ is reasonable, since conditioning can introduce new robustness problems. For example, suppose that a test statistic $Z_n(Y)$ tends in distribution to $N(0, 1)$ under H_0 as $n \rightarrow \infty$. In a non-selective setting, we might be comfortable modeling it as Gaussian as a basis for hypothesis testing. In this case it is also true that $\mathcal{L}(Z_n | Z_n > c)$ converges to a truncated Gaussian law for any fixed $c \in \mathbb{R}$, but the approximation may be much poorer for intermediate values of n . Worse, if we use increasing thresholds $c_n \rightarrow \infty$ with n , the truncated Gaussian approximation may never become reasonable.

2.6 Prior Work on Selective Inference

This article takes its main inspiration from a recent ferment of work on the problem of inference in linear regression models after model selection. Lockhart et al. (2014) derive an asymptotic test for whether the nonzero fitted coefficients at a given knot in the lasso path contain all of the true nonzero coefficients. Taylor et al. (2014) provided an exact (finite-sample) version of this result and extended it to the LARS path, while Lee et al. (2013), Loftus and Taylor (2014), and Lee and Taylor (2014) used similar approaches to derive exact tests for the lasso with a fixed value of regularization parameter λ , forward stepwise regression, and regression after marginal screening, respectively. All of the above approaches are derived assuming that the error variance σ^2 is known or an independent estimate is available.

The present work attempts to unify the above approaches under a common theoretical framework generalizing the classical optimality theory of Lehmann and Scheffé (1955), and elucidate previously unexplored questions of power. It also lets us generalize the results to the case of unknown σ^2 , and to arbitrary exponential families after arbitrary selection events.

Other works have viewed selective inference as a multiple inference problem. Recent work in this vein can be found in Berk et al. (2013) and Barber and Candès (2014). Section 8 argues that

inference after model selection and multiple inference are distinct problems with different scientific goals; see Benjamini (2010) for more discussion of this distinction. An empirical Bayes approach for selection-adjusted estimation can be found in Efron (2011).

There has also recently been work on inference in high-dimensional linear regression models, notably Belloni et al. (2011), Belloni et al. (2014), Zhang and Zhang (2014), Javanmard and Montanari (2013), and van de Geer et al. (2013); see Dezeure et al. (2014) for a review. These works focus on approximate asymptotic inference for a fixed model with many variables, while we consider finite-sample inference after selecting a smaller submodel to focus our inferential goals.

Leeb and Pötscher (2005, 2006, 2008) prove certain impossibility results regarding estimating the distribution of post-selection estimators. These results do not apply to our framework; under the statistical models we use, the post-selection distributions of our test statistics are known and thus do not require estimation.

The foregoing works are frequentist, as is this work. Because Bayesian inference conditions on the entire data set, conditioning first on a selection event typically has no operative effect on the posterior: if p and π are respectively the marginal likelihood and prior, then $p(Y | A, \theta) \cdot \pi(\theta | A) \propto p(Y | \theta) \cdot \pi(\theta)$ for $Y \in A$ (Dawid, 1994). Yekutieli (2012) argues that in certain cases it is more appropriate to use the likelihood conditional on selection condition the likelihood on selection without changing the prior to reflect that conditioning, resulting in a posterior proportional to $p(Y | A, \theta) \cdot \pi(\theta)$. The credible intervals discussed in Yekutieli (2012) resemble the confidence intervals proposed in this article, and the discussion therein presents a somewhat different perspective on how and why conditioning can adjust for selection.

Though our goals are very different, our theoretical framework is in some respects similar to the conditional confidence framework of Kiefer (1976), in which inference is made conditional on some estimate of the confidence with which a decision can be made. See also Kiefer (1977); Brownie and Kiefer (1977); Brown (1978); Berger et al. (1994).

Olshen (1973) discussed error control given selection in a two-stage multiple comparison procedure, in which an F -test is first performed, then Scheffé’s S -method applied if the F -test rejects. For large enough rejection thresholds, simultaneous coverage in the second stage is less than $1 - \alpha$ conditional on rejection in stage one.

3 Selective Inference in Exponential Families

As discussed in Section 2.2, we can construct selective tests “one at a time” for each model-hypothesis pair (M, H_0) , conditional on the corresponding selection event A_q and ignoring any other models that were previously under consideration. This is because the other candidate models and hypotheses are irrelevant to satisfying (9). For that reason, we suppress the explicit dependence on $q = (M, H_0)$ except where it is necessary to resolve ambiguity.

Our framework for selective inference is especially convenient when M corresponds to a multiparameter exponential family

$$Y \sim f_\theta(y) = \exp\{\theta' T(y) - \psi(\theta)\} f_0(y) \quad (22)$$

with respect to some dominating measure. Then, the conditional distribution given $Y \in A$ for any measurable A is another exponential family with the same natural parameters and sufficient statistics but different carrier distribution and normalizing constant:

$$(Y | Y \in A) \sim \exp\{\theta' T(y) - \psi_A(\theta)\} f_0(y) \mathbf{1}_A(y) \quad (23)$$

This fact lets us draw upon the rich theory of inference in multiparameter exponential families.

3.1 Conditional Inference and Nuisance Parameters

Classically, conditional inference in exponential families arises as a means for inference in the presence of nuisance parameters, as in Model 4 below.

Model 4 (Exponential Family with Nuisance Parameters). *Y follows a p -parameter exponential family with sufficient statistics $T(y)$ and $U(y)$, of dimension k and $p - k$ respectively:*

$$Y \sim f_{\theta, \zeta}(y) = \exp\{\theta' T(y) + \zeta' U(y) - \psi(\theta, \zeta)\} f_0(y), \quad (24)$$

with $(\theta, \zeta) \in \Theta \subseteq \mathbb{R}^p$ open.

Assume θ corresponds to a parameter of interest and ζ to an unknown nuisance parameter. The conditional law $\mathcal{L}(T(Y) | U(Y))$ depends only on θ :

$$(T | U = u) \sim g_\theta(t | u) = \exp\{\theta' t - \psi_g(\theta | u)\} g_0(t | u), \quad (25)$$

letting us eliminate ζ from the problem by conditioning on U . For $k = 1$ (i.e., for $\theta \in \mathbb{R}$), we obtain a single-parameter family for T .

Consider testing the null hypothesis $H_0 : \theta \in \Theta_0 \subseteq \Theta$ against the alternative $H_1 : \theta \in \Theta_1 = \Theta \setminus \Theta_0$. We say a level- α selective test $\phi(y)$ is *selectively unbiased* if

$$\text{Pow}_\phi(\theta | A) = \mathbb{E}_\theta[\phi(Y) | A] \geq \alpha, \quad \text{for all } \theta \in \Theta_1. \quad (26)$$

The condition (26) specializes to the usual definition of an unbiased test when there is no selection (when $A = \mathcal{Y}$). Unbiasedness rules out tests that privilege some alternatives to the detriment of others, such as one-sided tests when the alternative is two-sided.

A *uniformly most powerful unbiased* (UMPU) selective level- α test is one whose selective power is uniformly highest among all level- α tests satisfying (26). A selectively unbiased confidence region is one that inverts a selectively unbiased test, and confidence regions inverting UMPU selective tests are called uniformly most accurate unbiased (UMAUI). All of the above specialize to the usual definitions when $A = \mathcal{Y}$.

See Lehmann and Romano (2005) or Brown (1986) for thorough reviews of the rich literature on testing in exponential family models. In particular, the following classic result of Lehmann and Scheffé (1955) gives a simple construction of UMPU tests in exponential family models.

Theorem 5 (Lehmann and Scheffé (1955)). *Under Model 4 with $k = 1$, consider testing the hypothesis*

$$H_0 : \theta = \theta_0 \quad \text{against} \quad H_1 : \theta \neq \theta_0 \quad (27)$$

at level α . There is a UMPU test of the form $\phi(Y) = f(T(Y), U(Y))$ with

$$f(t, u) = \begin{cases} 1 & t < c_1(u) \text{ or } t > c_2(u) \\ \gamma_i & t = c_i(u) \\ 0 & c_1(u) < t < c_2(u) \end{cases} \quad (28)$$

where c_i and γ_i are chosen to satisfy

$$\mathbb{E}_{\theta_0}[f(T, U) | U = u] = \alpha \quad (29)$$

$$\mathbb{E}_{\theta_0}[T f(T, U) | U = u] = \alpha \mathbb{E}_{\theta_0}[T | U = u]. \quad (30)$$

The condition (29) constrains the power to be α at $\theta = \theta_0$, and (30) is obtained by differentiating the power function and setting its derivative to 0 at $\theta = \theta_0$.

Because $\mathcal{L}(Y | A)$ is an exponential family, we can simply apply Theorem 5 to the conditional law $\mathcal{L}(Y | A)$ to obtain an analogous construction in the selective setting.

Corollary 6 (UMPU Selective Tests). *Under Model 4 with $k = 1$, consider testing the hypothesis*

$$H_0 : \theta = \theta_0 \quad \text{against} \quad H_1 : \theta \neq \theta_0 \quad (31)$$

at selective level α on selection event A . There is a UMPU selective test of the form $\phi(Y) = f(T(Y), U(Y))$ with

$$f(t, u) = \begin{cases} 1 & t < c_1(u) \text{ or } t > c_2(u) \\ \gamma_i & t = c_i(u) \\ 0 & c_1(u) < t < c_2(u) \end{cases} \quad (32)$$

for which c_i and γ_i solve

$$\mathbb{E}_{\theta_0} [f(T, U) | U = u, Y \in A] = \alpha \quad (33)$$

$$\mathbb{E}_{\theta_0} [Tf(T, U) | U = u, Y \in A] = \alpha \mathbb{E}_{\theta_0} [T | U = u, Y \in A]. \quad (34)$$

It is worth keeping in mind that unbiasedness is only one way to choose a test when there is no completely UMP one. For example, another simple choice is to use the equal-tailed test from the same conditional law (25). The equal-tailed level- α rejection region is simply the union of the one-sided level- $\alpha/2$ rejection regions. While the equal-tailed and UMPU tests choose c_i and γ_i in different ways, both tests take the form (28). In fact, as we will see next, *all* admissible tests are of this form, which implies that data splitting tests are usually inadmissible.

3.2 Conditioning, Admissibility, and Data Splitting

A selective level- α test ϕ is *inadmissible* on selection event A if there exists another selective level- α test ϕ^* for which

$$\mathbb{E}_{\theta, \zeta} [\phi^*(Y) | A] \geq \mathbb{E}_{\theta, \zeta} [\phi(Y) | A], \quad \text{for all } (\theta, \zeta) \in \Theta_1, \quad (35)$$

with the inequality strict for at least one (θ, ζ) . In the main result of this section, we will show that tests based on data splitting are nearly always inadmissible.

Let Y be an observation from Model 4, and suppose we wish to test

$$H_0 : \theta = \theta_0 \quad \text{against} \quad H_1 : \theta \neq \theta_0. \quad (36)$$

We will assume all tests are functions of the sufficient statistic and write (with some abuse of notation) $\phi(T, U)$ for $\phi(Y)$. We can do this without loss of generality because any test $\phi(Y)$ can be Rao-Blackwellized, i.e.,

$$\phi(T, U) \equiv \mathbb{E}[\phi(Y) | T, U],$$

to obtain a new test that is a function of (T, U) , with the same power function as the original. Therefore, if $\phi(T, U)$ is inadmissible, then so is the original test $\phi(Y)$.

Now we can apply the following result of Matthes and Truax (1967).

Theorem 7 (Matthes and Truax, Theorem 3.1). *Let Y be an observation from Model 4, and suppose we wish to test*

$$H_0 : \theta = \theta_0 \quad \text{against} \quad H_1 : \theta \neq \theta_0. \quad (37)$$

Let \mathcal{C} denote the class of all level- α tests $\phi(T, U)$ of the form

$$\phi(t, u) = \begin{cases} 0 & t \in \text{int } C(u) \\ \gamma(t, u) & t \in \partial C(u) \\ 1 & t \notin C(u) \end{cases}, \quad (38)$$

and $C(u)$ is a convex set for every u . Then, for any $\phi \notin \mathcal{C}$, there exists $\phi^ \in \mathcal{C}$ such that*

$$\mathbb{E}_{\theta, \zeta} [\phi^*(T, U)] \geq \mathbb{E}_{\theta, \zeta} [\phi(T, U)], \quad \text{for all } (\theta, \zeta) \in \Theta_1. \quad (39)$$

Notice that, if (39) holds with equality for all (θ, ζ) , then by the completeness of (T, U) we have $\phi \stackrel{\text{a.s.}}{=} \phi^*$. Hence, every admissible test is in \mathcal{C} or almost surely equal to a test in \mathcal{C} .

In order to apply this result to data splitting, we first introduce a generic exponential family composed of two independent data sets governed by the same parameters:

Model 8 (Exponential Family with Data Splitting). *Model independent random variables $(Y_1, Y_2) \in \mathcal{Y}_1 \times \mathcal{Y}_2$ as*

$$Y_i \sim \exp \{ \theta T_i(y) + \zeta' U_i(y) - \psi_i(\theta, \zeta) \} f_{0,i}(y), \quad i = 1, 2, \quad (40)$$

with $\theta \in \mathbb{R}$ and with the models for Y_i both satisfying Model 4.

Model 8 would, for example, cover the case where Y_1 and Y_2 are the responses for two linear regressions with different design matrices but the same regression coefficients.

For a selection event $A = A_1 \times \mathcal{Y}_2$, we say ϕ is a *data-splitting test* if $\phi(Y) = \phi_2(Y_2)$; that is, the selection stage uses only Y_1 and the inference stage uses only Y_2 . Again, by Rao-Blackwellization, we can assume w.l.o.g. that the test is of the form $\phi(T_2, U_2)$.

Next, define the *cutoff gap* $g^*(\phi)$ as the largest $g \geq 0$ for which the acceptance and rejection regions are separated by a “cushion” of width g . If T_2^* is a conditionally independent copy of T_2 given U_2 , then

$$g^*(\phi) = \sup \{ g : \mathbb{P}_{\theta, \zeta}(|T_2 - T_2^*| < g, \phi(T_2, U_2) > 0, \phi(T_2^*, U_2) < 1) = 0 \}. \quad (41)$$

Note that the support of (T_2, T_2^*, U_2) does not depend on θ or ζ ; thus, neither does g^* . For most tests, $g^*(\phi) = 0$. For example, $g^* = 0$ if either cutoff is in the interior of $\text{supp}(T_2 | U_2)$ with positive probability, or if ϕ is a randomized test for discrete (T_2, U_2) .

Next we prove the main technical result of this section: ϕ is inadmissible unless T_1 is determined by U_1 on A_1 , within an amount g^* of variability.

Theorem 9. *Let T_1^* denote a copy of T_1 that is conditionally independent given U_1 and $Y_1 \in A$, and let ϕ be a data-splitting test of (37) in Model 8. If*

$$\mathbb{P}_{\theta, \zeta}(|T_1 - T_1^*| > g^*(\phi) | Y_1 \in A) > 0$$

then ϕ is inadmissible.

Proof. Construct conditionally independent copies T_i^* with $(T_1, T_1^*, U_1) \perp\!\!\!\perp (T_2, T_2^*, U_2)$, and assume that ϕ is of the form $\phi(T, U)$ with $T = T_1 + T_2$ and $U = U_1 + U_2$ (otherwise we could Rao-Blackwellize it). If ϕ is admissible, then by Matthes and Truax (1967), it must be a.s. equivalent to a test of the form (38). That is, there exist $c_i(U)$ for which

$$\mathbb{P}_{\theta, \zeta}(\phi(T, U) < 1, T \notin [c_1(U), c_2(U)] | A) = \mathbb{P}_{\theta, \zeta}(\phi(T, U) > 0, c_1(U) < T < c_2(U) | A) = 0. \quad (42)$$

Now, by assumption, there exists $\delta > g^*(\phi)$ for which

$$B_1 \triangleq \{|T_1 - T_1^*| > \delta\}$$

occurs with positive probability. By the definition of $g^*(\phi)$ in (41), the event

$$B_2 \triangleq \{|T_2 - T_2^*| > \delta, \phi(T_2, U_2) > 0, \phi(T_2^*, U_2) < 1\}$$

also occurs with positive probability. Since the two events are independent, $B = B_1 \cap B_2$ occurs with positive probability.

Next, assume w.l.o.g. that the event in (41) can occur with $T_2^* > T_2$ (otherwise we could reparameterize with natural parameter $\xi = -\theta$, for which $-T_i$ would be the sufficient statistics for Y_i). Then for some $\delta > g^*(\phi)$, the event

$$B = \{T_1 + \delta < T_1^*, T_2 < T_2^* < T_2 + \delta, \phi(T_2) > 0, \text{ and } \phi(T_2^*) < 1\}$$

occurs with positive probability for all θ, ζ . On B ,

$$T_1 + T_2 < T_1 + T_2^* < T_1^* + T_2 < T_1 + T_2^*,$$

but $\phi(T, U) > 0$ for $T = T_1 + T_2$ and $T = T_1^* + T_2$ and $\phi(T, U) < 1$ for the other two, ruling out the possibility of (42). \square

In the typical case $g^* = 0$ and we have

Corollary 10. *Suppose ϕ is a data-splitting test of (37) in Model 8 with $g^*(\phi) = 0$. Then ϕ is inadmissible unless T_1 is a function of U_1 (that is, unless $T_1 \in m\mathcal{F}(U_1)$).*

Example 3. To illustrate Theorem 9, consider a bivariate version of Example 2:

$$Y_i \sim N(\mu, 1), \quad i = 1, 2, \quad \text{with } Y_1 \perp\!\!\!\perp Y_2, \quad (43)$$

in which we condition on the selection event $A = \{Y_1 > 3\}$.

With data splitting, we could construct a 95% confidence interval using only Y_2 ; namely, $Y_2 \pm 1.96$. This interval is valid but does not use all the information available. A more powerful alternative is to construct an interval based on the law

$$\mathcal{L}_\mu(Y_1 + Y_2 \mid Y_1 > 3), \quad (44)$$

which uses the leftover information in Y_1 .

Figure 4a shows the Fisher information that is available to each test as a function of μ . The Fisher information of data splitting is exactly 1 no matter what μ is, whereas the optimal selective test has information approaching 2 as μ increases. Figure 4b shows the expected confidence interval length of the equal tailed interval as a function of μ . For $\mu \gg 3$, the data splitting interval is roughly 41% longer than it needs to be (in the limit, the factor is $\sqrt{2} - 1$).

Together, the plots tell a consistent story: when the selection event is not too unlikely, discarding the first data set exacts an unnecessary toll on the power of our second-stage procedure.

4 Selective Inference for Linear Regression

For a concrete example of the exponential family framework discussed in Section 3, we now turn to linear regression, which is one of the most important applications of selective inference. In linear regression, the data arise from a multivariate normal distribution

$$Y \sim N_n(\mu, \sigma^2 I_n),$$


where μ is modeled as

$$\mu = X_M \beta^M. \quad (45)$$

To avoid trivialities, we will assume that X_M has full column rank for all M under consideration, so that β^M is well-defined.

Depending on whether σ^2 is assumed known or unknown, hypothesis tests for coordinates β_j^M generalize either the z -test or the t -test. In the non-selective case, z - and t -tests are based on coordinates of the ordinary least squares (OLS) estimator $\hat{\beta} = X_M^\dagger Y$, where X_M^\dagger is the Moore-Penrose pseudoinverse. For a particular j and M , it will be convenient to write $\hat{\beta}_j^M = \eta_j^M Y$ with

$$\eta_j^M = \frac{X_{j \cdot M}}{\|X_{j \cdot M}\|^2}, \quad \text{where } X_{j \cdot M} = \mathcal{P}_{X_{M \setminus j}}^\perp X_j \quad (46)$$



(a) Fisher information available for second-stage inference.

(b) Expected confidence interval length.

Figure 4: Contrast between data splitting and data carving in Example 3, in which $Y_i \sim N(\mu, 1)$ independently for $i = 1, 2$. Data splitting discards Y_1 entirely, while data carving uses the leftover information in Y_1 for the second-stage inference. When $\mu \ll 3$, data carving also uses about one data point for inference since there is no information left over in Y_1 . But when $\mu \gg 3$, conditioning barely effects the law of Y_1 and data carving has nearly two data points left over.

is the remainder after adjusting X_j for the other columns of X_M , and $\mathcal{P}_{X_{M \setminus j}}$ denotes projection onto the column space of $X_{M \setminus j}$. Letting $\hat{\sigma}^2 = \|\mathcal{P}_{X_M}^\perp Y\|^2 / (n - |M|)$, the test statistics

$$Z = \frac{\eta_j^{M'} Y}{\sigma \|\eta_j^M\|} \quad \text{and} \quad \tilde{T} = \frac{\eta_j^{M'} Y}{\hat{\sigma} \|\eta_j^M\|} \quad (47)$$

are respectively distributed as $N(0, 1)$ and $t_{n-|M|}$ under $H_0 : \beta_j^M = 0$. Henceforth, we will suppress the subscript and superscript for η_j^M , simply writing η when there is no ambiguity.

We will see that the optimal selective versions of these tests are based on the same test statistics, but compared against different null distributions.

4.1 Inference Under the Selected Model

Suppressing the superscript M in β^M , the selected model has the form

$$Y \sim \exp \left\{ \frac{1}{\sigma^2} \beta' X_M' y - \frac{1}{2\sigma^2} \|y\|^2 - \psi(X_M \beta, \sigma^2) \right\} \quad (48)$$

If σ^2 is known, the sufficient statistics are $X_j' Y$ for $k \in M$, and inference for β_j is based on

$$\mathcal{L}_{\beta_j} (X_j' Y \mid X_{M \setminus j}' Y, A). \quad (49)$$

Otherwise, $\|Y\|^2$ represents another sufficient statistic and inference is based on

$$\mathcal{L}_{\beta_j / \sigma^2} (X_j' Y \mid X_{M \setminus j}' Y, \|Y\|, A). \quad (50)$$

Decomposing

$$X_j' y = X_j' \mathcal{P}_{X_{M \setminus j}} y + X_j' \mathcal{P}_{X_{M \setminus j}}^\perp y \quad (51)$$

$$= X_j' \mathcal{P}_{X_{M \setminus j}} y + \|X_{j \cdot M}\|^2 \eta' y, \quad (52)$$

we see that $Z = \eta'Y/\sigma\|\eta\|$ is a fixed affine transformation of $X_j'Y$ once we condition on $X_{M\setminus j}'Y$. If σ^2 is known, then, we can equivalently base our selective test on

$$\mathcal{L}_{\beta_j}(Z \mid X_{M\setminus j}'Y, A). \quad (53)$$

While Z is marginally independent of $X_{M\setminus j}'Y$, it is generically not conditionally independent given A , so that the null distribution of Z generically depends on $X_{M\setminus j}'Y$.

If σ^2 is unknown, we may observe further that

$$\hat{\sigma}^2 = \frac{\|\mathcal{P}_{X_M}^\perp Y\|^2}{n - |M|} = \frac{\|Y\|^2 - \|\mathcal{P}_{X_{M\setminus j}} Y\|^2 - (\eta'Y)^2/\|\eta\|^2}{n - |M|}. \quad (54)$$

Writing $Z_0(Y) = \eta'Y/\|\eta\|$, we have $\tilde{T}(Y) = (n - |M|) Z_0/(\|Y\|^2 - \|\mathcal{P}_{X_{M\setminus j}} Y\|^2 - Z_0^2)$, which is a monotone function of $\eta'Y$ after fixing $\|Y\|^2$ and $X_{M\setminus j}'Y$. Thus, our test is based on the law of

$$\mathcal{L}_{\beta_j/\sigma^2}(\tilde{T} \mid X_{M\setminus j}'Y, \|Y\|, A). \quad (55)$$

Note that, given A , $\hat{\sigma}^2$ in (54) is neither unbiased for σ^2 nor χ^2 -distributed. We recommend against viewing it as a serious estimate of σ^2 in the selective setting.

Constructing a selective t -interval is not as straightforward as the general case described in Section 5.2 because β_j is not a natural parameter of the selected model; rather, β_j/σ^2 is. Testing $\beta_j = 0$ is equivalent to testing $\beta_j/\sigma^2 = 0$, but testing $\beta_j = c$ for $c \neq 0$ does not correspond to any point null hypothesis about β_j/σ^2 . However, we can define

$$\tilde{Y} = Y - bX_j \sim N(X\beta - bX_j, \sigma^2 I). \quad (56)$$

Because $(\beta_j - b)/\sigma^2$ is a natural parameter for \tilde{Y} , we can carry out a UMPU selective t -test for $H_0 : \beta_j = b \iff (\beta_j - b)/\sigma^2 = 0$ based on the law of \tilde{Y} .

4.2 Inference Under the Saturated Model

Even if we do not take the linear model (45) seriously, there is still a well-defined best linear predictor in the population for design matrix X_M :

$$\theta^M = \arg \min_{\theta} \mathbb{E}_{\mu} [\|Y - X_M \theta\|^2] = X_M^\dagger \mu, \quad (57)$$

We call θ^M the *least squares coefficients* for M . According to this point of view, each θ_j^M corresponds to the linear functional $\eta_j^{M'} \mu$.

This point of view is convenient because the least-squares parameters are well-defined under the more general saturated model (6), leading to meaningful inference even if we do a poor job of selecting predictors. In particular, Berk et al. (2013) adopt this perspective as a way of avoiding the need to consider multiple candidate probabilistic models.

Several recent articles have tackled the problem of exact selective inference in linear regression after specific selection procedures (Lee et al., 2013; Loftus and Taylor, 2014; Lee and Taylor, 2014). These works, as well as Berk et al. (2013), assume the error variance is known, or that an estimate may be obtained from independent data, and target least-squares parameters in the saturated model.

Under the selected model, $\beta_j^M = \theta_j^M = \eta_j' \mu$, whereas under the saturated model β^M may not exist (i.e., there is no β^M such that $\mu = X_M \beta^M$). Compared to the selected model, the saturated model has $n - |M|$ additional nuisance parameters corresponding to $\mathcal{P}_{X_M}^\perp \mu$.

We can write the saturated model in exponential family form as

$$Y \sim \exp \left\{ \frac{1}{\sigma^2} \mu' y - \frac{1}{2\sigma^2} \|y\|^2 - \psi(\mu, \sigma^2) \right\}, \quad (58)$$

which has $n + 1$ natural parameters if σ^2 is unknown and n otherwise. To perform inference on some least-squares coefficient $\theta_j^M = \eta' \mu$, we can rewrite (58) as

$$Y \sim \exp \left\{ \frac{1}{\sigma^2 \|\eta\|^2} \mu' \eta \eta' y + \frac{1}{\sigma^2} (\mathcal{P}_\eta^\perp \mu)' (\mathcal{P}_\eta^\perp y) - \frac{1}{2\sigma^2} \|y\|^2 - \psi(\mu, \sigma^2) \right\}. \quad (59)$$

If σ^2 is known, inference for θ_j^M after selection event A is based on the conditional law $\mathcal{L}_{\theta_j^M}(\eta' Y \mid \mathcal{P}_\eta^\perp Y, A)$, or equivalently $\mathcal{L}_{\theta_j^M}(Z \mid \mathcal{P}_\eta^\perp Y, A)$.

If σ^2 is unknown, we must instead base inference on

$$\mathcal{L}_{\theta_j^M/\sigma^2}(\eta' Y \mid \mathcal{P}_\eta^\perp Y, \|Y\|, A). \quad (60)$$

Unfortunately, the conditioning in (60) is too restrictive. The set

$$\{y : \mathcal{P}_\eta^\perp y = \mathcal{P}_\eta^\perp Y, \|y\| = \|Y\|\} \quad (61)$$

is a line intersected with the sphere $\|Y\|S^{n-1}$, and consists only of the two points $\{Y, Y - 2\eta'Y\}$, which are equally likely under the hypothesis $\theta_j^M = 0$. Thus, under the saturated model, conditioning on $\|Y\|$ leaves insufficient information about θ_j^M to carry out a meaningful test.

4.3 Saturated Model or Selected Model?

When σ^2 is known, we have a choice whether to carry out the z -test with test statistic $Z = \eta'Y/\sigma\|\eta\|$ in the saturated or the selected model. In other words, we must choose either to assume that $\mathcal{P}_{X_M}^\perp \mu = 0$ or to treat it as an unknown nuisance parameter. Writing

$$U = X_{M \setminus j}' Y, \quad \text{and} \quad V = \mathcal{P}_{X_M}^\perp Y, \quad (62)$$

we must choose whether to condition on U and V (saturated model) or only U (selected model). Conditioning on both U and V can never increase our power relative to conditioning only on U , and will in most cases lead to an inadmissible test per Theorem 9.


In the non-selective case, this choice makes no difference at all since T, U , and V are mutually independent. In the selective case, however, the choice may be of major consequence as it can lead to very different tests. In general, T, U , and V are not conditionally independent given A , and $\mathcal{P}_{X_M}^\perp \mu$ may play an important role in determining the conditional distribution of T . If we needlessly condition on V , we may lose a great deal of power, whereas failing to condition on V could lead us astray if $\mathcal{P}_{X_M}^\perp \mu$ is large. A simple example can elucidate this contrast.

Example 4. Suppose that $y \sim N_2(\mu, I_2)$, with design matrix $X = I_2$, and we choose the best one-sparse model. Our selection procedure chooses $M = \{1\}$ if $|Y_1| > |Y_2|$, and $M = \{2\}$ otherwise.

Figure 5 shows one realization of this process with $Y = (2.9, 2.5)$. $|Y_1|$ is a little larger than $|Y_2|$, so we choose $M = \{1\}$. The yellow highlighted region $A = \{|Y_1| > |Y_2|\}$ is the chosen selection event, and the selected model is

$$Y \sim N_2((\mu_1, 0), I_2). \quad (63)$$

In this case, $T = Y_1$, $V = Y_2$, and there is no U since X_M has only one column. The selected-model test is based on $\mathcal{L}(Y_1 \mid A)$, whereas the saturated-model test is based on $\mathcal{L}(Y_1 \mid Y_2, A)$. The second conditioning set, a union of two rays, is plotted in brown. Under the hypothesis $\mu = 0$, the realized $|Y_1|$ is quite large given A , giving p -value 0.015. By contrast, $|Y_1|$ is not terribly large given $\{Y_2 = 2.5\} \cap A = \{Y_2 = 2.5, |Y_1| > 2.5\}$, leading to p -value 0.30.



(a) For $Y = (2.9, 2.5)$, the selected-model conditioning set is $A = \{y : |y_1| > |y_2|\}$, a union of quadrants, plotted in yellow. The saturated-model conditioning set is $\{y : y_2 = 2.5\} \cap A = \{y : y_2 = 2.5, |y_1| > 2.5\}$, a union of rays, plotted in brown.

(b) Conditional distributions of Y_1 under $H_0 : \mu_1 = 0$. Under the hypothesis $\mu = 0$, the realized $|Y_1|$ is quite large given A , giving p -value 0.015. By contrast, $|Y_1|$ is not too large given $A \cap \{y : y_2 = Y_2\}$, giving p -value 0.3.

Figure 5: Contrast between the saturated-model and selected-model tests in Example 4, in which we fit a one-sparse model with design matrix $X = I_2$. The selected-model test is based on $\mathcal{L}_0(Y_1 | A)$, whereas the saturated-model test is based on $\mathcal{L}_0(Y_1 | Y_2, A)$.

The difference between the saturated and selected models is especially important in early steps of sequential model-selection procedures that use a form of the saturated-model z -test. It has been observed in several cases that if there are two strong variables with similar effect sizes, the p -value in the first step may not be very small (Taylor et al., 2014; G'Sell et al., 2013). We will explore this subtle issue further in a forthcoming companion paper dealing with sequential model selection.

5 Computations

We saw in Section 3 that inference in the one-parameter exponential family requires knowing the conditional law $\mathcal{L}_\theta(T | U, A)$. In a few cases, such as in the saturated model viewpoint, this conditional law can be determined fairly explicitly. In other cases, we will need to resort to Monte Carlo sampling. In this section, we suggest some general strategies.

5.1 Gaussians Under the Saturated Model

As we discussed in Section 4.2, the previous papers by Lee et al. (2013); Loftus and Taylor (2014); Lee and Taylor (2014) adopted the saturated model viewpoint with known σ^2 . In this case, $\mathcal{L}_\theta(T | U, A) = \mathcal{L}_\theta(\eta'Y \mid \mathcal{P}_\eta^\perp Y, A)$ is a truncated univariate Gaussian, since $\eta'Y$ is a Gaussian random variable and $\mathcal{P}_\eta^\perp Y$ is independent of $\eta'Y$. If A is convex, then the truncation is to an interval $[\mathcal{V}^-(Y), \mathcal{V}^+(Y)]$, where the endpoints represent the maximal extent one can move in the




Figure 6: Saturated-model inference for a generic convex selection set for $Y \sim N(\mu, I_n)$. After conditioning on the yellow set A , \mathcal{V}^+ is the largest $\eta'Y$ can get while \mathcal{V}^- is the smallest it can get. Under $H_0 : \eta'\mu = 0$, the test statistic $\eta'Y$ takes on the distribution of a standard Gaussian random variable truncated to the interval $[\mathcal{V}^-, \mathcal{V}^+]$. As a result, $W(Y) = \frac{\Phi(\eta'Y) - \Phi(\mathcal{V}^-)}{\Phi(\mathcal{V}^+) - \Phi(\mathcal{V}^-)}$ is uniformly distributed.

η direction at a “height” of $\mathcal{P}_\eta^\perp Y$, while still remaining inside A , i.e.,

$$\mathcal{V}^+(Y) = \sup_{\{t: Y+t\eta \in A\}} \eta'(Y+t\eta) \quad (64)$$

$$\mathcal{V}^-(Y) = \inf_{\{t: Y+t\eta \in A\}} \eta'(Y+t\eta). \quad (65)$$

The geometric intuition is illustrated in Figure 6.

When A is specifically a polytope, we can obtain closed-form expressions for \mathcal{V}^- and \mathcal{V}^+ . The generalization to regions A that are non-convex is straightforward (i.e., instead of truncating to a single interval, we truncate to a union of intervals). For further discussion of these points, see Lee et al. (2013).

5.2 Monte Carlo Tests and Intervals

More generally, if we can obtain a stream of samples from $\mathcal{L}_\theta(T|U, A)$ for any value of θ , then we can carry out hypothesis tests and construct intervals. This can be done efficiently via rejection sampling if, for example, we can sample efficiently from $\mathcal{L}_\theta(Y|U)$ and $\mathbb{P}_\theta(Y \in A|U)$ is not too small. Otherwise, more specialized sampling approaches may be required.

More generally, we consider how to construct a test based on the statistic Z , which is distributed according to a one-parameter exponential family

$$Z \sim g_\theta(z) = e^{\theta z - \psi(\theta)} g_0(z). \quad (66)$$

If in addition to Z we are given an independent sequence from the reference distribution

$$Z_1, \dots, Z_n \stackrel{\text{i.i.d.}}{\sim} g_0(z), \quad (67)$$

then an exact Monte Carlo one-sided test of $H_0 : \theta \leq 0$ rejects if the observed value Z is among the $(n+1)\alpha$ largest of Z, Z_1, \dots, Z_n (Barnard, 1963).

By reweighting the samples, we can use the same sequence to test $H_0 : \theta \leq \theta_0$ for any other θ_0 . Denote the importance-weighted empirical expectation as

$$\hat{\mathbb{E}}_\theta h(Z) = \frac{\sum_{i=1}^n h(Z_i) e^{\theta Z_i}}{\sum_{i=1}^n e^{\theta Z_i}} \quad (68)$$

$$\xrightarrow{a.s.} \mathbb{E}_\theta h(Z) \quad \text{as } n \rightarrow \infty \text{ for integrable } h. \quad (69)$$

In effect, we have put an exponential family “through” the empirical distribution of the Z_i in the manner of Efron et al. (1996); see also Besag (2001). The Monte Carlo one-sided cutoff for a test of $H_0 : \theta \leq \theta_0$ is the smallest c_2 for which

$$\hat{\mathbb{P}}_{\theta_0}(Z > c_2) \leq \alpha. \quad (70)$$

The test rejects for $Z > c_2$ and randomizes appropriately at $Z = c_2$.

The Monte Carlo UMPU two-sided test of $H_0 : \theta = \theta_0$ is a bit more involved, but similar in principle. We can solve for $c_1, \gamma_1, c_2, \gamma_2$ for which

$$\hat{\mathbb{E}}_{\theta_0} \phi(Z) = \alpha \quad (71)$$

$$\hat{\mathbb{E}}_{\theta_0} [Z \phi(Z)] = \alpha \hat{\mathbb{E}}_{\theta_0} Z. \quad (72)$$

In Appendix B we discuss how (71–72) can be solved efficiently for fixed θ_0 and inverted to obtain a confidence interval. Monte Carlo inference as described above is computationally straightforward once Z_1, \dots, Z_n are obtained.

More generally, the Z_i could represent importance samples with weights W_i , or steps in a Markov chain with stationary distribution $g_{\theta_0}(z)$. The same methods apply as long as we still have

$$\hat{\mathbb{E}}_\theta h(Z) = \frac{\sum_{i=1}^n W_i h(Z_i) e^{\theta Z_i}}{\sum_{i=1}^n W_i e^{\theta Z_i}} \quad (73)$$

$$\xrightarrow{a.s.} \mathbb{E}_\theta h(Z), \quad \text{for integrable } h. \quad (74)$$

Numerical problems may arise in solving (71–72) for θ_0 far away from the reference parameter used for sampling. Combining appropriately weighted samples from several reference values θ can help to keep the effective sample size from getting too small for any θ_0 . For further references on Monte Carlo inference see Jockel (1986); Besag and Clifford (1989); Forster et al. (1996); Mehta et al. (2000).

5.3 Sampling Gaussians with Affine and Quadratic Constraints

In the case where Y is Gaussian, several simplifications are possible. For one, there are many ways to sample from a truncated multivariate Gaussian distribution. In this paper, we use hit-and-run Gibbs sampling algorithms, while Pakman and Paninski (2014) suggest another approach based on Hamiltonian Monte Carlo.

Efficient sampling from multivariate Gaussian distributions under such constraints is the main algorithmic challenge for most of the Gaussian selective tests proposed in this paper. The works cited above use the saturated model exclusively which means they do not require any sampling.

In many cases, the sampling problem may be greatly facilitated by refining the selection variable that we use. For example, Lee et al. (2013) propose conditioning on the variables selected by the lasso as well as the signs of the fitted $\hat{\beta}_j$, leading to a selection event consisting of a single polytope in \mathbb{R}^n . If we condition only on the selected variables and not on the signs, the selection event is a union of up to 2^s polytopes, where s is the number of variables in the selected model (though most of the polytopes might be excluded after conditioning on U).

Refining the selection variable never impairs the selective validity of the procedure, but it typically leads to a loss in power. However, this loss of power may be quite small if, for example, the conditional law puts nearly all of its mass on the realized polytope. This price in power is acceptable if it is the only way to obtain a tractable test. Quantifying the tradeoff between computation and power is an interesting topic for further work.

When carrying out selective t -tests, it is necessary to condition further on the realized vector length $\|Y\|$, adding a quadratic equality constraint to the support. To deal with this, we sample instead from a ball and project the samples onto the sphere using an importance sampling scheme. Appendix C gives details.

6 Selective Inference in Non-Gaussian Settings

In this section we describe tests in two simple non-Gaussian settings, selective inference in a binomial problem, and tests involving a scan statistic in Poisson process models. More generally, we address the question of selective inference in generalized linear models.

6.1 Selective Clinical Trial

To illustrate the application of our approach in a simple non-Gaussian setting we discuss a selective clinical trial with binomial data. The experiment discussed here is similar to an adaptive design proposed by Sill and Sampson (2009).

Consider a clinical trial with m candidate treatments for heart disease. We give treatment j to n_j patients for $0 \leq j \leq m$, with $j = 0$ corresponding to the placebo. The number of patients on treatment j to suffer a heart attack during the trial is

$$Y_j \stackrel{\text{ind.}}{\sim} \text{Binom}(p_j, n_j), \quad \text{with } \log \frac{p_j}{1-p_j} = \begin{cases} \theta & j = 0 \\ \theta - \beta_j & j > 0 \end{cases}, \quad (75)$$

so β_j measures the efficacy of treatment j . The likelihood for Y is

$$Y \sim \exp \left\{ \theta \sum_{j=0}^m y_j - \sum_{j=1}^m \beta_j y_j - \psi(\theta, \beta) \right\} \prod_{j=0}^m \binom{Y_j}{n_j}, \quad (76)$$

an exponential family with $m+1$ sufficient statistics. Define $\hat{p}_j = Y_j/n_j$, and let $\hat{p}_{(j)}$ denote the j th smallest order statistic.

After observing the data, we select the best $k < m$ treatments in-sample, then construct a confidence interval for each one's odds ratio relative to placebo. If there are ties, we select all treatments for which $\hat{p}_j \leq \hat{p}_{(k)}$ (so that we could possibly select more than k treatments).

For simplicity, assume that treatments $1, \dots, k$ are the ones selected. Inference for β_1 is then based on the conditional law

$$\mathcal{L}_{\beta_1} \left(Y_1 \mid \sum_{j=0}^m Y_j, Y_2, \dots, Y_m, \{j = 1 \text{ selected}\} \right) \quad (77)$$

Under this law, Y_2, \dots, Y_m are fixed, as is $Y_0 + Y_1$, with Y_0 and Y_1 the only remaining unknowns. Before conditioning on selection, we have the two-by-two multinomial table

	Control	Treatment
Heart attack	Y_0	Y_1
No heart attack	$n_0 - Y_0$	$n_1 - Y_1$

The margins are fixed, and conditioning on selection gives an additional constraint that $Y_1 \leq n_1 \hat{p}_{(k)}$, where the right-hand side is known after conditioning on the other Y_j . Rejecting for conditionally extreme Y_1 amounts to a selective Fisher's exact test. Aside from the constraint on its support, the distribution of Y_1 is hypergeometric if $\beta_1 = 0$ and otherwise noncentral hypergeometric with noncentrality parameter β_1 . We can use this family to construct an interval for β_1 .

6.2 Poisson Scan Statistic

As a second simple example, consider observing a Poisson process $Y = \{Y_1, \dots, Y_{N(Y)}\}$ on the interval $[0, 1]$ with piecewise-constant intensity, possibly elevated in some unknown window $[a, b]$. That is, $Y \sim \text{Poisson}(\lambda(t))$ with

$$\lambda(t) = \begin{cases} e^{\alpha+\beta} & t \in [a, b] \\ e^{\alpha} & \text{otherwise.} \end{cases} \quad (78)$$

Our goal is to locate $[a, b]$ by maximizing some scan statistic, then test whether $\beta > 0$ or construct a confidence interval for it. Assume we always have $[\hat{a}, \hat{b}] = [Y_i, Y_j]$ for some i, j ; this is true, for example, if we use the multi-scale-adjusted likelihood ratio statistic proposed in Rivera and Walther (2013).

The density of Y can be written in exponential family form as

$$Y \sim \exp \left\{ \sum_{i=1}^{N(y)} \log \lambda(y_i) - \int_0^1 \lambda(s) ds \right\} \quad (79)$$

$$= \exp \{ \alpha N(Y) + \beta T(y) - \psi(\alpha, \beta) \}, \quad (80)$$

where

$$T(y) = \sum_{i=1}^{N(y)} \mathbf{1}\{y_i \in [a, b]\} \quad \text{and} \quad \psi(\alpha, \beta) = e^{\alpha}(1 - b + a) + e^{\alpha+\beta}(b - a). \quad (81)$$

If A is the event that $[a, b]$ is chosen, we carry out inference with respect to $\mathcal{L}_{\beta}(T | N, A)$. Note that under $\beta = 0$ and conditional on N , Y is an i.i.d. uniform random sample on $[0, 1]$.

Once we condition on the event $\{a, b \in Y\}$, the other $N - 2$ values are uniform. Thus, we can sample from $\mathcal{L}_{\beta}(T | N, A)$ with $\beta = 0$ by taking Y to include a, b , and $N - 2$ uniformly random points, then rejecting samples for which $[a, b]$ is not the selected window.

6.3 Generalized Linear Models

Our framework extends to logistic regression, Poisson regression, or other generalized linear model (GLM) with response Y and design matrix X , since the GLM model may be represented as an exponential family of the form

$$Y \sim \exp \{ \beta' X' y - \psi(X\beta) \} f_0(y). \quad (82)$$

As a result, we can proceed just as we did in the case of linear regression in the reduced model, conditioning on $U = X_{M \setminus j}' Y$ and basing inference on $\mathcal{L}_{\beta_j^M}(X_j' Y | U, A)$.

A difficulty may arise for logistic or Poisson regression due to the discreteness of the response distribution Y . If some control variable X_1 is continuous, then for almost every realization of X , all configurations of Y yield unique values of $U = X_1' Y$. In that case, conditioning on $X_1' Y$ means conditioning on Y itself. No information is left over for inference, so that the best (and only) exact

level- α selective test is the trivial one $\phi(Y) \equiv \alpha$. By contrast, if all of the control variables are discrete variables like gender or ethnicity, then conditioning on U may not constrain Y too much.

Because $X'Y$ is approximately a multivariate Gaussian random variable, a more promising approach may to base inference on the asymptotic Gaussian approximation, though we will not pursue that here. Tian and Taylor (2015) discuss selective inference in certain non-Gaussian problems.

7 Simulation: High-Dimensional Regression

As a simple illustration, we compare selective inference in linear regression after the lasso for $n = 100, p = 200$. Here, the rows of the design matrix X are drawn from an equicorrelated multivariate Gaussian distribution with pairwise correlation $\rho = 0.3$ between the variables. The columns are normalized to have length 1.

We simulate from the model

$$Y \sim N(X\beta, I_n), \quad (83)$$

with β 7-sparse and its non-zero entries set to 7. The signal to noise ratio (SNR) (magnitude of β) was chosen so that data splitting with half the data yielded a superset of the true variables on roughly 20% of instances. For data splitting and carving, Y is partitioned into selection and inference data sets Y_1 and Y_2 , containing n_1 and $n_2 = n - n_1$ data points respectively.

We compare two procedures:

Data Splitting after Lasso on Y_1 (Split $_{n_1}$): Use the lasso on Y_1 to select the model, and use Y_2 for inference.

Data Carving after Lasso on Y_1 (Carve $_{n_1}$): Use the lasso on Y_1 to select the model, and use Y_2 and whatever is left over of Y_1 for inference.

Procedure Carve $_{100}$ is inference after the using the lasso on the full data set Y .

For the data carving procedures, we use the selected-model z -test of Section 4.1 after lasso variable selection using Lagrange parameter

$$\lambda = 2\mathbb{E}(\|X^T \epsilon\|_\infty), \quad \epsilon \sim N(0, \sigma^2 I)$$

as described in (Negahban et al., 2012). In addition, we condition on the signs of the active lasso coefficients, so that procedure Carve $_{100}$ is the inference-after-lasso test considered in Lee et al. (2013).⁴

We know from Theorem 9 that procedure Carve $_{n_1}$ strictly dominates procedure Split $_{n_1}$ for any n_1 , but there is a selection–inference tradeoff between data-carving procedures Carve $_n$ and Carve $_{n_1}$ for $n_1 < n$. Carve $_n$ uses all of the data for selection, and is therefore likely to select a superior model, whereas procedure Carve $_{n_1}$ reserves more power for the second stage.

Let R be the size of the model selected and V the number of noise variables included. We compare the procedures with respect to aspects of their selection performance:

- chance of screening, i.e. obtaining a correct model ($\mathbb{P}(R - V = 7)$ or p_{screen}).
- expected number of noise variables selected ($\mathbb{E}[V]$),
- expected number of true variables selected ($\mathbb{E}[R - V]$),

⁴Because of the form of the selection event when we use the lasso after n data points, the test statistic is conditionally independent of $\mathcal{P}_{X_M}^\perp Y$. Thus, there is no distinction between the saturated- and selected-model z -tests after the lasso on all n data points.

Algorithm	p_{screen}	$\mathbb{E}[V]$	$\mathbb{E}[R - V]$	FDR	Power	Level
Carve ₁₀₀	0.99	8.13	6.99	0.54	0.80	0.05
Split ₅₀	0.09	9.13	4.74	0.66	0.93	0.06
Carve ₅₀	0.09	9.13	4.74	0.66	0.99	0.06
Split ₇₅	0.68	9.24	6.59	0.58	0.47	0.05
Carve ₇₅	0.68	9.24	6.59	0.58	0.97	0.06

Table 1: Simulation results. p_{screen} is the probability of successfully selecting all 7 true variables, and Power is the power, conditional on successful screening, of tests on the true variables. The more data we use for selection, the better the selected model’s quality is, but there is a cost in second-stage power. Carve₇₅ appears to be finding a good tradeoff between these competing goals. Carve _{n_1} always outperforms Split _{n_1} , as predicted by Theorem 9.

Algorithm	p_{screen}	$\mathbb{E}[V]$	$\mathbb{E}[R - V]$	FDR	Power	Level
Carve ₁₀₀	0.97	8.11	6.97	0.54	0.80	0.04
Split ₅₀	0.09	9.20	4.77	0.66	0.93	0.05
Carve ₅₀	0.09	9.20	4.77	0.66	0.99	0.06

Table 2: Simulation results under misspecification. Here, errors ϵ are drawn independently from Student’s t_5 . Our conclusions are identical to Table 1.


- false discovery rate of true variables selected ($\mathbb{E}[V/\max(R, 1)]$ or FDR),

Conditional on having obtained a correct model, we also compare them on aspects of their second stage performance:

- probability of correctly rejecting the null for one of the true variables (Power),
- probability of incorrectly rejecting the null for a noise variable (Level).

The results, shown in Table 1, bear out the intuition of Section 3.2. Because procedure Carve₁₀₀ uses the most information in the first stage, it performs best in terms of model selection, but pays a price in lower second-stage power relative to Split₅₀ or Carve₅₀. The procedure Carve₅₀ clearly dominates Split₅₀, as expected. Increasing n_1 from 50 to 75 improves p_{screen} for Split₇₅, but Split₇₅ suffers a drop in power. Procedure Carve₇₅ seems to strike a better compromise.

Figure 7 shows the tradeoff curve of model selection success (as measured by the probability of successful screening) against second-stage power conditional on successful screening. As n_1 increases, stage-one performance improves while stage-two performance declines, but the decline is much slower for data carving. Surprisingly, Carve₉₈ and Carve₉₉ have much higher power than Carve₁₀₀: 91%, 86%, and 80% respectively. We cannot explain why holding out just one or two data points in the first stage improves power so dramatically. Better understanding this tradeoff is an interesting topic of further work.



(a) Probability of successful screening, and power conditional on screening, for Split_{n_1} and Carve_{n_1} . (b) Probability of successful screening times power conditional on screening, for Split_{n_1} and Carve_{n_1} .

Figure 7: Tradeoff between power and model selection. As n_1 increases and more data is used in the first stage, we have a better chance of successful screening (picking all the true nonzero variables). However, increasing n_1 also leads to reduced power in the second stage. Data splitting suffers much more than data carving, though both are affected.

Finally, to check the robustness of data carving, we replace the Gaussian errors with independent errors drawn from Student's t distribution with five degrees of freedom. The numbers barely change at all; see Table 2. Tian and Taylor (2015) rigorously analyze the case of non-Gaussian errors.

8 Selective Inference and Multiple Inference

A common approach to the problem of inference after selection is to replace the nominal error rate with an alternative joint error rate that is deemed appropriate to the scientific setting.

For example, suppose that θ_q , $q = 1, \dots, m$ correspond to parameters of a common model M . We designate a small number $R(Y) = |\hat{\mathcal{Q}}(Y)|$ of them as interesting and construct a confidence interval $C_q(Y)$ for each $q \in \hat{\mathcal{Q}}$. Benjamini and Yekutieli (2005) propose controlling the *false coverage-statement rate* (FCR)

$$\mathbb{E} \left[\frac{V}{\max(R, 1)} \right], \quad \text{where} \quad V(Y) = \left| \left\{ q : q \in \hat{\mathcal{Q}}, \theta_q(F) \notin C_q(Y) \right\} \right| \quad (84)$$

is the number of non-covering intervals constructed.

FCR control is closely related to selective coverage. By choosing appropriate selection variables S_q , we can adapt selective coverage to control the FCR.

Proposition 11 (FCR Control via Selective Error Control). *Assume \mathcal{Q} is countable with each $q \in \mathcal{Q}$ corresponding to a different parameter θ_q for the same model M . Let $R(Y) = |\hat{\mathcal{Q}}(Y)|$ with $R(Y) < \infty$ a.s., and define $V(Y)$ as in (84).*

If each C_q enjoys coverage at level $1 - \alpha$ given $S_q = (\mathbf{1}_{A_q}(Y), R(Y))$, then the collection of intervals $(C_q, q \in \hat{\mathcal{Q}})$ controls the FCR at level α :

$$\mathbb{E} \left[\frac{V}{\max(R, 1)} \right] \leq \mathbb{E} \left[\frac{V}{R} \mid R \geq 1 \right] \leq \alpha. \quad (85)$$

Proof. Let $V_q(Y) = \mathbf{1}\{q \in \hat{Q}(Y), \theta_q(F) \notin C_q(Y)\}$, so that $V = \sum_{q \in \mathcal{Q}} V_q$. For $R \geq 1$, and for any $F \in M$,

$$\mathbb{E}_F[V | R] = \sum_{q \in \mathcal{Q}} \mathbb{E}_F[V_q | R] \leq \sum_{q \in \mathcal{Q}} \alpha \mathbb{E}_F[\mathbf{1}_{A_q}(Y) | R] = \alpha R, \quad (86)$$

hence $\mathbb{E}[V/R | R] = \alpha$ for each $R \geq 1$. \square

Other authors have addressed inference after selection by proposing to control the FWER, the chance that any selected test incorrectly rejects the null or any constructed confidence interval fails to cover its parameter. For example, the “post-selection inference” (PoSI) method of Berk et al. (2013) constructs simultaneous $(1 - \alpha)$ confidence intervals for the least-squares parameters of all linear regression models that were ever under consideration. As a result, no matter how we choose the model, the overall probability of constructing any non-covering interval is controlled at α .

Selective error control can be adapted to control the FWER as well. If our selection rule always chooses a single hypothesis to test, then we have overall FWER control.

Proposition 12 (FWER Control for Singleton \hat{Q}). *Assume that $|\hat{Q}(Y)| \stackrel{a.s.}{=} 1$, and let $Q(Y) = (M(Y), H_0(Y))$ denote the single (random) selected model-hypothesis pair. If each ϕ_q controls the selective error at level α , then the test $\phi(y) = \phi_{Q(y)}(y)$ controls the FWER at level α :*

$$\mathbb{E}_F[\phi(Y); F \in H_0(Y)] \leq \alpha \quad (87)$$

Proof. Condition on Q :

$$\mathbb{E}_F[\phi(Y)\mathbf{1}\{F \in H_0(Y)\}] = \mathbb{E}_F[\mathbb{E}_F[\phi(Y) | Q(Y)] \mathbf{1}\{F \in H_0(Y)\}] \quad (88)$$

$$\leq \alpha \mathbb{P}_F[F \in H_0(Y)] \quad (89)$$

\square

More generally, it is clear that if we construct $(\phi_q, q \in \mathcal{Q})$ to control any joint error rate conditional on the entire selected set \hat{Q} , we will also have marginal control of the same joint error rate.

However, the converse of Proposition 12 is not true: FWER control does *not* in general guarantee control of relevant selective error rates. For example, suppose $Q(Y) = 1$ with probability 0.9 and $Q(Y) = 2$ otherwise. If ϕ_1 and ϕ_2 have selective error rates $\alpha_1 = 0.02$ and $\alpha_2 = 0.3$ respectively, the overall FWER is still controlled at $\alpha = 0.05$.

Does our conservatism when asking question 1 compensate for our anti-conservatism when asking question 2? To answer this question we must consider not only mathematics but also the relevant scientific context. If the different questions represent a bag of anonymous, *a priori* undifferentiated hypotheses, then a joint error rate like the FWER or FDR may be a good proxy for our scientific goals. For example, when performing a large-scale genome-wide “fishing expedition” for loci associated with type II diabetes, the fraction of null genes among all purported discoveries is a very relevant quantity: it measures what fraction of our time and money will be wasted following up on false leads.

In other scientific applications, however, different hypotheses have quite distinct identities and may vary greatly in their importance and interpretation. For example, a confidence interval for the effect of gender on salary after controlling for one’s job title may be much more socially consequential than an interval for the effect of job title after controlling for gender. As such, averaging our error rates across the two questions is inappropriate.

9 Discussion

Selective inference concerns the properties of inference carried out after using a data-dependent procedure to select which questions to ask. We can recover the same long-run frequency properties among answers to *selected* questions that we would obtain in the classical non-adaptive setting, if we follow the guiding principle of selective error control:

The answer must be valid, given that the question was asked.

Happily, living up to this principle can be a simple matter in exponential family models including linear regression, due to the rich classical theory of optimal testing in exponential family models. Even if we are possibly selecting from a large menu of diverse and incompatible models, we can still design tests one model at a time and control the selective error using the test designed for the selected model. We generally pay a price for conditioning, so it is desirable to condition on as little as possible. Data carving can dramatically improve on data splitting by using the leftover information in Y_1 , the data set initially designated for selection.

Many challenges remain. Deriving the cutoffs for sample carving tests can be computationally difficult in general. In addition, the entire development of this article takes the model selection procedure \hat{Q} as given, when in reality we can choose \hat{Q} . More work is needed to learn what model selection procedures lead to favorable second-stage properties.

As data sets and research questions become more and more complex, we have less and less hope of specifying adequate probabilistic models ahead of time. As such, a key challenge of complex research is to balance the goal of choosing a realistic model against the goal of inference once we have chosen it. We hope that the ideas in this article represent a step in the right direction.

Reproducibility

A git repository with code to generate the figures for this file is available at the first author's website.

Acknowledgements

William Fithian was supported by National Science Foundation VIGRE grant DMS-0502385 and the Gerald J. Lieberman Fellowship. Dennis Sun was supported in part by the Stanford Genome Training Program (NIH/NHGRI T32 HG000044) and the Ric Weiland Graduate Fellowship. Jonathan Taylor was supported in part by National Science Foundation grant DMS-1208857 and Air Force Office of Sponsored Research grant 113039. We would like to thank Stefan Wager, Trevor Hastie, Rob Tibshirani, Brad Efron, Yoav Benjamini, Larry Brown, Maxwell Grazier G'sell, Subhabrata Sen, and Yuval Benjamini for helpful discussions.

References

- Rina Foygel Barber and Emmanuel Candes. Controlling the false discovery rate via knockoffs. *arXiv preprint arXiv:1404.5609*, 2014.
- GA Barnard. Discussion of professor bartlett's paper. *Journal of the Royal Statistical Society*, 1963.
- Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. Inference for high-dimensional sparse econometric models. *arXiv preprint arXiv:1201.0220*, 2011.

- Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650, 2014.
- Yoav Benjamini. Simultaneous and selective inference: current successes and future challenges. *Biometrical Journal*, 52(6):708–721, 2010.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300, 1995.
- Yoav Benjamini and Daniel Yekutieli. False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469):71–81, 2005.
- James O Berger, Lawrence D Brown, and Robert L Wolpert. A unified conditional frequentist and bayesian test for fixed and sequential simple hypothesis testing. *The Annals of Statistics*, pages 1787–1807, 1994.
- Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, and Linda Zhao. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- Julian Besag. Markov chain monte carlo for statistical inference. *Center for Statistics and the Social Sciences*, 2001.
- Julian Besag and Peter Clifford. Generalized monte carlo significance tests. *Biometrika*, 76(4): 633–642, 1989.
- Lawrence D Brown. A contribution to kiefer’s theory of conditional confidence procedures. *The Annals of Statistics*, pages 59–71, 1978.
- Lawrence D Brown. Fundamentals of statistical exponential families with applications in statistical decision theory. *Lecture Notes-monograph series*, pages i–279, 1986.
- C Brownie and J Kiefer. The ideas of conditional confidence in the simplest setting. *Communications in Statistics-Theory and Methods*, 6(8):691–751, 1977.
- Arthur Cohen and Harold B Sackrowitz. Two stage conditionally unbiased estimators of the selected mean. *Statistics & Probability Letters*, 8(3):273–278, 1989.
- DR Cox. A note on data-splitting for the evaluation of significance levels. *Biometrika*, 62(2): 441–444, 1975.
- AP Dawid. Selection paradoxes of bayesian inference. *Lecture Notes-Monograph Series*, pages 211–220, 1994.
- Ruben Dezeure, Peter Bühlmann, Lukas Meier, and Nicolai Meinshausen. High-dimensional inference: Confidence intervals, p-values and r-software hdi. *arXiv preprint arXiv:1408.4026*, 2014.
- Bradley Efron. Tweedies formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.
- Bradley Efron, Robert Tibshirani, et al. Using specially designed exponential families for density estimation. *The Annals of Statistics*, 24(6):2431–2461, 1996.
- Jonathan J Forster, John W McDonald, and Peter WF Smith. Monte carlo exact conditional tests for log-linear and logistic models. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 445–453, 1996.

- Annie Franco, Neil Malhotra, and Gabor Simonovits. Publication bias in the social sciences: unlocking the file drawer. *Science*, 2014.
- Andrew Gelman and Eric Loken. The garden of forking paths: Why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time. *Downloaded January*, 30:2014, 2013.
- Max Grazier G’Sell, Jonathan Taylor, and Robert Tibshirani. Adaptive testing for the graphical lasso. *arXiv preprint arXiv:1307.4765*, 2013.
- Naftali Harris. Visualizing lasso polytope geometry, June 2014. URL <http://www.naftaliharris.com/blog/lasso-polytope-geometry/>.
- Trevor Hastie, Robert Tibshirani, Jerome Friedman, T Hastie, J Friedman, and R Tibshirani. *The elements of statistical learning*, volume 2. Springer, 2009.
- Larry V Hedges. Estimation of effect size under nonrandom sampling: The effects of censoring studies yielding statistically insignificant mean differences. *Journal of Educational and Behavioral Statistics*, 9(1):61–85, 1984.
- Larry V Hedges. Modeling publication selection effects in meta-analysis. *Statistical Science*, pages 246–255, 1992.
- John PA Ioannidis. Why most published research findings are false. *PLoS medicine*, 2(8):e124, 2005.
- Adel Javanmard and Andrea Montanari. Hypothesis testing in high-dimensional regression under the gaussian random design model: Asymptotic theory. *arXiv preprint arXiv:1301.4240*, 2013.
- Karl-Heinz Jockel. Finite sample properties and asymptotic efficiency of monte carlo tests. *The annals of Statistics*, pages 336–347, 1986.
- George Johnson. New truths that only one can see. *The New York Times*, 2014.
- Jack Kiefer. Admissibility of conditional confidence procedures. *The Annals of Statistics*, pages 836–865, 1976.
- Jack Kiefer. Conditional confidence statements and confidence estimators. *Journal of the American Statistical Association*, 72(360a):789–808, 1977.
- Jason D Lee and Jonathan E Taylor. Exact post model selection inference for marginal screening. *arXiv preprint arXiv:1402.5596*, 2014.
- Jason D Lee, Dennis L Sun, Yuekai Sun, and Jonathan E Taylor. Exact post-selection inference with the lasso. *arXiv preprint arXiv:1311.6238*, 2013.
- Hannes Leeb and Benedikt M Pötscher. Model selection and inference: Facts and fiction. *Econometric Theory*, 21(01):21–59, 2005.
- Hannes Leeb and Benedikt M Pötscher. Can one estimate the conditional distribution of post-model-selection estimators? *The Annals of Statistics*, pages 2554–2591, 2006.
- Hannes Leeb and Benedikt M Pötscher. Can one estimate the unconditional distribution of post-model-selection estimators? *Econometric Theory*, 24(02):338–376, 2008.
- EL Lehmann and Joseph P Romano. *Testing statistical hypotheses*. New York:. Springer, 2005.

- EL Lehmann and Henry Scheffé. Completeness, similar regions, and unbiased estimation: Part ii. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 15(3):219–236, 1955.
- Richard Lockhart, Jonathan Taylor, Ryan J Tibshirani, and Robert Tibshirani. A significance test for the lasso (with discussion). *The Annals of Statistics*, 42(2):413–468, 2014.
- Joshua R Loftus and Jonathan E Taylor. A significance test for forward stepwise model selection. *arXiv preprint arXiv:1405.3920*, 2014.
- Ted K Matthes and Donald R Truax. Tests of composite hypotheses for the multivariate exponential family. *The Annals of Mathematical Statistics*, pages 681–697, 1967.
- Cyrus R Mehta, Nitin R Patel, and Pralay Senchaudhuri. Efficient monte carlo methods for conditional logistic regression. *Journal of The American Statistical Association*, 95(449):99–108, 2000.
- Sahand N. Negahban, Pradeep Ravikumar, Martin J. Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of MM-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, November 2012. ISSN 0883-4237. doi: 10.1214/12-STS400. URL <http://projecteuclid.org/euclid.ss/1356098555>.
- Richard A Olshen. The conditional level of the ftest. *Journal of the American Statistical Association*, 68(343):692–698, 1973.
- Ari Pakman and Liam Paninski. Exact hamiltonian monte carlo for truncated multivariate gaussians. *Journal of Computational and Graphical Statistics*, 23(2):518–542, 2014.
- Camilo Rivera and Guenther Walther. Optimal detection of a jump in the intensity of a poisson process or in a density with likelihood ratio statistics. *Scandinavian Journal of Statistics*, 40(4):752–769, 2013.
- JD Rosenblatt and Yoav Benjamini. Selective correlations; not voodoo. *NeuroImage*, 103:401–410, 2014.
- Allan R Sampson and Michael W Sill. Drop-the-losers design: Normal case. *Biometrical Journal*, 47(3):257–268, 2005.
- Michael W Sill and Allan R Sampson. Drop-the-losers design: Binomial case. *Computational statistics & data analysis*, 53(3):586–595, 2009.
- Robert Sladek, Ghislain Rocheleau, Johan Rung, Christian Dina, Lishuang Shen, David Serre, Philippe Boutin, Daniel Vincent, Alexandre Belisle, Samy Hadjadj, et al. A genome-wide association study identifies novel risk loci for type 2 diabetes. *Nature*, 445(7130):881–885, 2007.
- Jonathan Taylor, Richard Lockhart, Ryan J Tibshirani, and Robert Tibshirani. Post-selection adaptive inference for least angle regression and the lasso. *arXiv preprint arXiv:1401.3889*, 2014.
- Xiaoying Tian and Jonathan Taylor. Asymptotics of selective inference. *arXiv preprint arXiv:1501.03588*, 2015.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- Sara van de Geer, Peter Bühlmann, Ya’acov Ritov, and Ruben Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *arXiv preprint arXiv:1303.0518*, 2013.

Asaf Weinstein, William Fithian, and Yoav Benjamini. Selection adjusted confidence intervals with more power to determine the sign. *Journal of the American Statistical Association*, 108(501):165–176, 2013.

Daniel Yekutieli. Adjusted bayesian inference for selected parameters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):515–541, 2012.

Ed Yong. Replication studies: Bad copy. *Nature*, 485(7398):298–300, 2012.

Cun-Hui Zhang and Stephanie S Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014.

Hua Zhong and Ross L Prentice. Bias-reduced estimators and confidence intervals for odds ratios in genome-wide association studies. *Biostatistics*, 9(4):621–634, 2008.

Sebastian Zöllner and Jonathan K Pritchard. Overcoming the winners curse: estimating penetrance parameters from case-control data. *The American Journal of Human Genetics*, 80(4):605–615, 2007.

A Proof of Proposition 1

Proof. For group i , let R_i be the number of true nulls selected, i.e.,

$$R_i = \left| \left\{ (M, H_0) : (M, H_0) \in \widehat{\mathcal{Q}}_i(Y_i), F_i \in H_0 \subseteq M \right\} \right|,$$

and let V_i denote the number of false rejections. If $Z_n^V = \sum_{i=1}^n V_i$ and $Z_n^R = \sum_{i=1}^n R_i$, then we need to show $\limsup_{n \rightarrow \infty} Z_n^V / Z_n^R \leq \alpha$.

By design, $0 \leq V_i \leq R_i$ and $\mathbb{E}(V_i) \leq \alpha \mathbb{E}(R_i)$. As a result, $\mathbb{E}[Z_n^V] / \mathbb{E}[Z_n^R] \leq \alpha$ for every n , so we just need to show that the two sums are not far from their expectations. Because

$$\sum_{i=1}^{\infty} \frac{\text{Var}(R_i)}{i^2} \leq B \sum_{i=1}^{\infty} \frac{1}{i^2} < \infty,$$

we can apply Kolmogorov’s strong law of large numbers to the independent but non-identical sequence R_1, R_2, \dots to obtain

$$\frac{1}{n} (Z_n^R - \mathbb{E} Z_n^R) \xrightarrow{a.s.} 0, \quad \text{so} \quad \left| \frac{Z_n^R}{\mathbb{E} Z_n^R} - 1 \right| \leq \left| \frac{\delta}{n} (Z_n^R - \mathbb{E} Z_n^R) \right| \xrightarrow{a.s.} 0.$$

As for Z_n^V , we have

$$\frac{1}{n} (Z_n^V - \mathbb{E} Z_n^V) \xrightarrow{a.s.} 0, \quad \text{so} \quad \frac{Z_n^V}{\mathbb{E} Z_n^R} - \alpha \leq \frac{\delta}{n} (Z_n^V - \mathbb{E} Z_n^V) \xrightarrow{a.s.} 0;$$

in other words, $Z_n^R / \mathbb{E} Z_n^R \xrightarrow{a.s.} 1$ and $\limsup_n Z_n^V / \mathbb{E} Z_n^R \stackrel{a.s.}{\leq} \alpha$. □

B Monte Carlo Tests and Confidence Intervals: Details

Assume Z arises from a one-parameter exponential family

$$Z \sim g_\theta(z) = e^{\theta z - \psi(\theta)} g_0(z). \quad (90)$$

We wish to compute (by Monte Carlo) the UMPU two-sided rejection region for the hypothesis $H_0 : \theta = \theta_0$. Let $U \sim \text{Unif}[0, 1]$ be an auxiliary randomization variable.

Define the dictionary ordering on $[0, 1]$:

$$(z_1, u_1) \prec (z_2, u_2) \iff z_1 < z_2 \text{ or } (z_1 = z_2 \text{ and } u_1 < u_2). \quad (91)$$

If $\Gamma_1 = (c_1, \gamma_1)$ and $\Gamma_2 = (c_2, 1 - \gamma_2)$, then the region

$$R_{\Gamma_1, \Gamma_2} = \{(z, u) : (z, u) \prec \Gamma_1 \text{ or } (z, u) \succ \Gamma_2\} \quad (92)$$

implements the rejection region for the test with cutoffs c_1, c_2 and boundary randomization parameters γ_1, γ_2 .

For $\Gamma_1 \prec \Gamma_2$, write

$$K_1(\Gamma_1, \Gamma_2; \theta) = \mathbb{P}_\theta(R_{\Gamma_1, \Gamma_2}) - \alpha \quad (93)$$

$$K_2(\Gamma_1, \Gamma_2; \theta) = \mathbb{E}_\theta(Z | (Z, U) \in R_{\Gamma_1, \Gamma_2}^C) - \mathbb{E}_\theta(Z), \quad (94)$$

so that the correct cutoffs Γ_i are those for which $K_1(\Gamma_1, \Gamma_2; \theta) = K_2(\Gamma_1, \Gamma_2; \theta) = 0$. For fixed θ , K_1 is decreasing in Γ_1 and increasing in Γ_2 , while K_2 is increasing in both Γ_1 and Γ_2 .

Let $(Z_1, W_1), (Z_2, W_2), \dots$ be a sequence of random variables for which

$$\hat{\mathbb{E}}_\theta^n h(Z) = \frac{\sum_{i=1}^n W_i h(Z_i) e^{\theta Z_i}}{\sum_{i=1}^n W_i e^{\theta Z_i}} \quad (95)$$

$$\xrightarrow{a.s.} \mathbb{E}_\theta h(Z). \quad (96)$$

for all integrable h . This would be true if (Z_i, W_i) are a valid i.i.d. sample or i.i.d. importance sample from g_0 , or if they come from a valid Markov Chain Monte Carlo algorithm.

If \hat{K}_i^n are defined analogously to K_i for $i = 1, 2$, with \mathbb{E}_θ and \mathbb{P}_θ replaced with their importance-weighted empirical versions $\hat{\mathbb{E}}_\theta^n$ and $\hat{\mathbb{P}}_\theta^n$, then $\hat{K}_i^n \xrightarrow{a.s.} K$ pointwise as $n \rightarrow \infty$, and \hat{K}_i^n satisfy the same monotonicity properties almost surely for each n . As a result, we have almost sure convergence on compacta for $(\hat{K}_1^n, \hat{K}_2^n)$:

$$\sup_{(\Gamma_1, \Gamma_2) \in G} \max_i \left\| \hat{K}_i^n(\Gamma_1, \Gamma_2; \theta) - K_i(\Gamma_1, \Gamma_2; \theta) \right\| \quad (97)$$

for each θ , for compact $G \in (\mathbb{R} \times [0, 1])^2$.

We carry out our tests by solving for Γ_1 and Γ_2 which solve \hat{K}_1^n and \hat{K}_2^n , in effect defining the UMPU tests for a one-parameter exponential family through the approximating empirical measure. Specifically, we can define

$$\hat{\Gamma}_2(\Gamma_1; \theta) = \inf \left\{ \Gamma_2 : \hat{K}_1^n(\Gamma_1, \Gamma_2; \theta) = 0 \right\}, \quad (98)$$

with $\hat{\Gamma}_2 = \infty$ if the set is empty. That is, for a given lower cutoff we define the upper cutoff to obtain a level- α acceptance region if that is possible. Then, $\hat{K}_2^n(\Gamma_1, \hat{\Gamma}_2(\Gamma_1; \theta); \theta)$ is an increasing function and we can solve it using binary search. Let \hat{R}_θ denote the rejection region so obtained.

Note that (z, u) is in the left-tail of \hat{R}_θ if and only if $\hat{K}_2^n((z, u), \hat{\Gamma}_2((z, u); \theta); \theta) < 0$. This fact, paired with an analogous test for whether (z, u) is in the right tail, gives us a quick way to carry out the test. It also allows us to quickly find the upper and lower confidence bounds for the approximating empirical family, via binary search.

C Sampling for the Selective t -Test: Details

Let $C \subseteq \mathbb{R}^k$ denote a set with nonempty interior and consider the problem of integrating some integrable function $h(y)$ against the uniform probability measure on $C \cap S^{k-1}$, where S^{k-1} is the unit sphere of dimension $k-1$, assuming the intersection is non-empty. Assume we are given an i.i.d. sequence of uniform samples Y_1, Y_2, \dots from $C \cap B^k$, where B^k is the unit ball.

Let $R \sim \frac{r^{k-1}}{k}$, so that if $Z \sim \text{Unif}(S^{k-1})$, then $Y = RZ \sim \text{Unif}(B^k)$. Let

$$W(Z) = \left(\int_0^1 \mathbf{1}\{rZ \in C\} \frac{r^{k-1}}{k} dr \right)^{-1} \quad (99)$$

We can use the Y_i for which $Z_i = Y_i/\|Y_i\| \in C$ as a sequence of importance samples with weights $W(Z_i)$, since

$$\mathbb{E}(h(Z) \mathbf{1}\{Y, Z \in C\} W(Z)) \quad (100)$$

$$= \int_{S^{k-1}} \int_0^1 h(z) \mathbf{1}\{z, rz \in C\} W(z) \frac{r^{k-1}}{k} dr dz \quad (101)$$

$$= \int_{S^{k-1}} h(z) \mathbf{1}\{z \in C\} dz \quad (102)$$

$$= \mathbb{E}(h(Z) \mathbf{1}\{Z \in C\}). \quad (103)$$

To carry out the selective t -test of $H_0 : \beta_j = 0$, we need to sample from

$$\mathcal{L}(\eta'Y \mid \mathcal{P}_{X_{M \setminus j}}Y, \|Y\|, A). \quad (104)$$

Let $U = \mathcal{P}_{X_{M \setminus j}}Y$, and let $Q \in \mathbb{R}^{n \times (n-|M|-1)}$ be such that $QQ' = \mathcal{P}_{X_{M \setminus j}}^\perp$. Then $L^2 \triangleq \|Q'Y\|^2 = \|Y\|^2 - \|U\|^2$ is fixed under the selection event. Let

$$C = \{v : U + Qv \in A\}, \quad (105)$$

so that $A_U = U + QC$, an $(n-|M|-1)$ -dimensional hyperplane intersected with A , is the event we would sample from for the selective z -test.

Under H_0 , Y is uniformly distributed on

$$(U + QC) \cap \|Y\|S^{n-1} = U + Q \left(C \cap LS^{n-|M|-2} \right). \quad (106)$$

Assume we can resample Y^* uniformly from $A_U \cap (U + LB^{n-|M|-1})$, which is just sampling from A_U with an additional quadratic constraint. Then $V^* = Q'(Y^* - U)$ is a sample from the ball of radius L , intersected with C . We can turn V^* into an importance-weighted sample from the sphere via the scheme outlined above; then, the same importance weight suffices to turn Y^* into a sample from the selective t -test conditioning set.